


The Empathetic Algorithm Leveraging AI for Next-Level CX



When she asked him to dance AI was shy.

Robert Dew

Published by CapFeather Global

Level 12, 31 Queen Street. Melbourne, VIC 3000 Australia

First edition 2023

Copyright © 2023 Capfeather Pty Ltd

Printed in Australia.

Reprints and permission service

Contact: office@capfeather.global

ISBN: 978-0-6458539-0-2 [hard cover]

ISBN: 978-0-6458539-1-9 [eBook]

All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other non-commercial uses permitted by copyright law.

Any opinions expressed in the chapters are those of the author. The information contained within is for informational purposes only and is not intended to serve as a source of definitive advice on any of the topics discussed. This book is presented as-is, without warranty of any kind, either express or implied. Reasonable measures were taken to ensure the accuracy and completeness of the information contained in this book. However, the information is not guaranteed to be error-free, current, or complete. The author and publisher are not responsible for any errors or omissions or for the results obtained from the use of such information.

Further the author and publisher will not be liable for any direct, indirect, incidental, special, consequential, or punitive damages, including but not limited to, damages for loss of profits, goodwill, use, data, or other intangible losses, arising out of or in connection with the use or performance of the information in this book.

Acknowledgements

This journey of writing has been one I could not have taken alone. Throughout the process, I have been fortunate to have the support, guidance, and companionship of many incredible individuals.

First and foremost, I am deeply indebted to my partners at CapFeather Global, Cyrus Allen and Lachlan Austin. Their entrepreneurial spirit and shared adventure in starting our boutique consulting firm in the customer strategy space have been nothing short of inspirational. Our relentless dedication to our mission has undoubtedly been a significant driver in my desire to pen this book.

Chau Phan, whose meticulous eyes combed through every chapter, has been integral to this project. Her commitment to ensuring every reference was in place and every sentence was just right is greatly appreciated. This book is as much a product of her diligent work as it is of my thoughts.

A special thanks to my wife, Sarah. Her introduction to ChatGPT and Dall-e has opened up a world of possibilities, leading to our fascinating collaboration with AI in the writing of this book. Her support and encouragement throughout the process have been unwavering.

OpenAI deserves immense recognition for leading the charge in democratizing the tools of Artificial Intelligence. Their pioneering work has not only made this book possible but is also reshaping the future of numerous fields, including customer strategy consulting.

Lastly, I would like to acknowledge the extraordinary AI companion I had in this journey, ChatGPT 4.0. What was initially an exploration into AI-assisted writing turned into a profound collaboration, drastically reducing the time it took to complete this book. I look forward to more such collaborations and witnessing the positive impact this tool can have on various fields.

--Robert Dew

Foreword

As our world confronts a near-term future shaped by the development and application of ever-more advanced technologies, it is critical that executives don't lose focus on the importance of customer experience in remaining competitive and relevant. It is similarly critical that they avoid the trap of applying technology simply to reduce cost, as tempting as that might be.

The transformative power of AI technologies is already radically altering business landscapes. The Covid pandemic not only led to a seismic shift towards digital technologies, but also underscored the crucial role of artificial intelligence (AI) in healthcare, retail, education, and a myriad of other sectors.

These developments, rapid and relentless, call for a systematic understanding of AI's role in the Customer Experience (CX). This book explores not just the application of AI to automation but delves into the softer aspect of 'empathy', often overlooked yet vital for impactful CX. This book is a comprehensive guide for businesses navigating this new normal, where digitization has ceased to be an option and become a necessity. It offers insights into the mind of the customer and the evolution of CX in the AI era. From leveraging AI for cost reduction to creating value growth, it provides a holistic understanding of the interplay between AI, empathic understanding, and CX.

Whether you're an entrepreneur adapting to the shifts in customer behaviour or a seasoned executive reimagining business strategy, this book is a tool to leverage AI for a deeper connection with your customers. As AI evolves and reshapes industries, this book can be your compass in the pursuit of excellent customer experiences.

--Cyrus Allen

Managing Partner

CapFeather Global

Contents

Acknowledgements	iii
Foreword.....	iv
Tables and Figures.....	viii
I. Introduction to a fast changing world.....	9
• Automation and CX.....	10
Rise of Empathic Machines	13
Topics Overview	14
• CX as a Strategy	15
• Personalization Drives CX Quality.....	16
• Recent AI Developments	18
• Artificial Empathy	22
II. Theory of Mind for Mind Reading.....	25
• Human Interactions	26
Verbal and Non-verbal Communication	28
• Artificial Theories of Mind	31
• CX Personas	35
• AI and Human Centred Design	37
III. CX Personalization.....	45
• Simple, Satisfying or Surprising	46
• Simple CX.....	48
• Satisfying CX	51
• Surprising CX.....	56
• Measuring CX.....	59
• Consistent Surprise.....	61
Varian’s Law and CX Strategy	62
Surprise at the Process Level	63
IV. Digitization for Cost Reduction.....	65
• Approaches to cost reduction.....	66
• Automation.....	70
• Past Chatbot Deployments	71


- More or Less Jobs? 74
- CX Lens on Innovation Execution..... 76
- V. Revenue Growth with Empathic AI 78
 - AI Driven CX for Growth 81
 - Awareness 84
 - Consideration 84
 - Decision 87
 - Purchase 88
 - Use..... 88
 - Support..... 90
 - Repeat and Refer 93
- VI. Finding AI Opportunities for CX 95
 - Pareto Analysis 96
 - Pareto in Practice 97
 - Variation is Significant 102
 - Theory of Constraints 107
 - Drum-Buffer-Rope 108
 - The Offer You Can't Refuse 109
 - Mafia Offers in Healthcare CX 111
 - Building the Road Map 115
- VII. The New Sources of Competitive Advantage..... 117
 - Resource-Based View of Strategy 119
 - Open Source Issues 124
 - Dynamic Capability View 125
 - Firm Evolution 127
 - Niche Evolution 128
 - Business Model Evolution 129
 - Sustaining AI Advantages..... 131
 - Data Differentiation..... 131
 - Customisation and Integration..... 134
 - Domain Expertise 135

- VIII. Trust in AI is Complex..... 136
 - Privacy..... 137
 - Effectiveness..... 140
 - Transparency..... 141
 - Implications..... 141
 - Building Trust..... 142
 - Business Trust..... 144
 - Technology Trust..... 144
 - Trusting AI..... 145
 - AI Trust Cohorts..... 147
 - Increasing AI Trust..... 148
 - Damaging AI Trust..... 150
 - Critical AI Trust Issues..... 154
- IX. Future Outlook:..... 158
 - Trend Analysis..... 159
 - ChatGPT and Varian’s Rule..... 162
 - Potential Scenarios..... 169
 - More Powerful Hardware on Offer..... 169
 - Quantum Computing is a Flop..... 171
 - Better Healthcare Saves 1% of GDP..... 173
 - Chatbots Get Even Smarter..... 175
 - Chatbots are Everywhere..... 177
 - User Interfaces Become More Human..... 178
 - Humans Become Cyborgs..... 180
- X. Conclusion:..... 182
- Index..... 188
- References..... 191

Tables and Figures

- FIGURE 1: CAPFUL and Personalisation Objectives.....47
- FIGURE 2: Population in Agriculture from 1400 to 2020 70
- FIGURE 3: Customer Journey Stages to Purchase 82
- FIGURE 4: Customer Journey Stages After Purchase 83
- FIGURE 5: AI in Customer Service 90
- FIGURE 6: Marketing’s 60/20 Pareto Law 98
- FIGURE 7: Pareto Chart of WeChat Activity 100
- FIGURE 8: Pareto Distribution of Consumer Packaged Goods 100
- FIGURE 9: Pareto Distribution of Department Store Customers..... 102
- FIGURE 10: Snapchat’s History 129
- FIGURE 11: Most Significant CX Trends Supported by AI..... 159
- FIGURE 12: Waves of AI Supported CX 161

I. Introduction to a fast changing world



AI couldn't believe another AI rejected his credit card.

Customer experience (CX) is changing because of how artificial intelligence (AI) is evolving. This book covers the latest AI developments and explores what they mean for CX strategy. AI is both an opportunity and a threat for companies because it radically repositions the potential of automating customer interactions. Early indications suggest that next generation open-source AI tools can reduce costs to serve by 90% and costs to analyse will reduce by 99%. Quality improvements for customer interactions are much harder to quantify. After a proof-of-concept demonstration of a next level AI based conversational agent, one potential investor described the chatbot as 'freakishly good'. The level of quality possible from automated interaction is now orders of magnitude above merely incremental improvement. Comparing last generation automated software to the current AI powered chatbots is like comparing TNT to nuclear weapons or pagers to mobile phones. This means executives and directors will have to come to terms with unprecedented market disruption.

Rethinking CX strategy is only part of the challenge. It is an important part because CX can bridge the gap between business as usual and innovation. Innovation is hard for established firms because they have learned to scale by reducing risk. All innovation suffers from the liability of newness because costs are front loaded and results are uncertain. This means innovation tends to activate the corporate immune system designed to protect the organisation from risks. Large firms tend to have an overactive corporate immune system. This is where CX strategy can help. Improving the customer experience the firm delivers is innovating, even if it doesn't seem to be risky.

Recent advancements in large language models are revolutionizing the way businesses can approach CX, paving the way for a new era of personalized interactions that are both efficient and empathetic. Prior to the current developments in large language models, companies had limited choices to deliver customer experiences at touch points. They could use people, who were more expensive but offered higher quality responsiveness and emotional connection, or automated systems, which were less expensive but offered lower quality responsiveness and no emotional connection. Increasingly companies have been choosing automated systems and rationalizing their labour force to capture cost savings at the expense of deeper relationships with their customers. Many companies have divested labour in favour of chatbots wherever possible. The problem has been chatbots are not as good as humans for customers (until recently).

- Automation and CX

Chatbots, a common form of automated systems, lack the ability to manage exceptional circumstances that require human judgment and going the extra mile for the customer. This inability leads to lower perceived value premium

and a decreased willingness to pay for the service, compared to when humans are involved. For instance, consider the perception of food quality with different levels of automation. Food from a vending machine like Febo in the Netherlands is typically perceived as lower quality than food purchased at a McDonald's drive-thru, where there is more human interaction. However, the drive-thru is also typically perceived as a lower quality food service than that in a sit-down restaurant, where humans take your order and personally prepare your food. In the limit, having a personal chef who asks you what you want to eat and goes beyond your request to surprise you with a twist is perceived as highest quality food option. In fact, this is no longer perceived as meal, rather as a 'dining experience.' In all these cases the nutritional value of the food is not the main determining factor that drives perceived quality. What drives perceived quality is the degree humans are involved in the customer journey to choose, order, produce and deliver the food. The emotional connection and responsiveness of humans leads to a higher value premium and increased willingness to pay as compared to dealing with an automated system. This applies in many different vertical markets and is true for both B2C and B2B purchase situations.

A 2014 blog post on TopRank Marketing emphasized the importance of the human element in B2B marketing and highlights that finding common ground and shared goals can unite different departments, tactics, and technologies¹. It asserted "The B2B industry will be re-crowned BP2BP, or business person-to-business person. However, not everyone is prepared to acknowledge this 'empathy' premium is foregone with automation. Perhaps a decade later times have changed.

According to McKinsey research, B2B buyers are increasingly comfortable with making large purchases through digital self-serve channels². In their study, they found that 70% of B2B decision-makers were open to making new, fully self-serve or remote purchases more than \$50,000, and 27% would spend more than \$500,000³. McKinsey asserts this shift in purchasing behaviour highlights the growing acceptance of digital sales channels among B2B buyers and emphasizes the importance of providing seamless online experiences to meet their needs. We believe this research is flawed methodologically.

Intention to purchase surveys are commonly used in market research to predict customer behaviour and measure the likelihood of a customer making a purchase. However, these surveys come with certain inaccuracies and risks that need to be considered when interpreting the results. Some of these inaccuracies and risks include inaccurate data, self-generated validity effects, sampling error, non-response error, and perceived risk. In short, some people who state they intend to buy don't and others who state they do not intend to buy go on and purchase. Perhaps McKinsey should have better supervised the juniors they got to prepare their report.

Other research indicates that a more human approach in B2B marketing and sales can lead to better customer experiences, stronger relationships, and increased business success. A study by Deloitte Digital shows that B2B companies that excel in customer experience see buyers being 34% more likely to purchase and 32% more likely to renew contracts with them⁴. An Allison+Partners report reveals that 97% of B2B brands want to become more human and focus on building human connections and engaging conversations with their audiences⁵. However, they often struggle to implement this approach effectively. Note these two papers are subject to the same methodology flaws as the McKinsey research because they also used intention to purchase surveys.

Thus, while the McKinsey survey may suggest that some B2B firms are open to high-value purchases via online systems, the reality is that human interaction is still the preferred method for many. Any B2B firms open to purchasing high-value products online would expect to receive a substantial discount, due to the economics of automation. If lower cost is the driving factor in their interest, this suggests the strategy is the starting gun on another 'race to the bottom.' Strategies designed to win with lower pricing can be rational if based on cost leadership, but such advantages rarely last in contemporary markets. Despite this, a lot of money has been invested in automation in the last decade.

The limitations of machines in providing an emotional connection and making judgment calls have led to the development of recommendation engines as a first order solution. The problem remains that these engines still fall short in understanding the nuanced preferences of customers, leaving human interaction as the preferred method for high-value purchases in B2B firms. Recommendation engines are more apparent and relevant in the B2C space.

Systems designed to prompt additional purchases have been widely deployed in B2C contexts. These engines use algorithms to analyse a customer's past behaviour and make recommendations based on what other consumers with similar behaviour have done. A company that has successfully used recommendation engines is Netflix⁶. Their engine analyses a customer's viewing history and recommends movies and TV shows based on their past behaviour. Another company that has rapidly grown because of its use of recommendation engines is Spotify, which uses their engine to suggest songs and playlists based on a user's listening history⁷. Both examples fail to prompt additional purchases however because consumers typically deal with each platform on an 'all you can eat' subscription basis. In contrast, Stitch Fix, an online personal styling service, uses a recommendation engine to suggest clothing items for their customers based on their personal style preferences and previous purchases⁸. Customers complete a style profile and provide feedback on the items they receive in each shipment. The recommendation

engine analyses this data and suggests new items for future shipments. The engine considers a range of factors such as customer preferences, purchase history, and feedback to provide personalized recommendations. The customer then has the option to purchase the items suggested to them.

Stitch Fix could deliver even better CX if they combined the recommendation engine with the empathy of human operators as a kind of expert system. This is becoming more common in the beauty industry. Many beauty brands use recommendation engines to suggest products to their customers based on their previous purchases and preferences. However, these engines alone cannot understand the specific needs and concerns of the customer's skin or hair type, nor can they provide personalized advice or emotional connection.

To address this, some beauty brands have combined recommendation engines with human expertise to create a hybrid system. For example, Sephora, a major beauty retailer, has a program called "Sephora Virtual Artist" that uses augmented reality technology to allow customers to try on different makeup looks⁹. The program also has a feature called "Sephora Color IQ" that uses a handheld device to scan a customer's skin and then recommends foundation shades that match their skin tone.

However, to provide additional support and expertise, Sephora also employs beauty advisors who are available to answer customers' questions, provide personalized recommendations, and offer emotional connection. By combining the recommendation engine with human expertise, Sephora provides a more personalized and satisfying customer experience. This hybrid approach has been best practice in CX up until the recent development of Large Language Models (LLMs).

Rise of Empathic Machines

LLMs have the potential to go beyond hybrid systems where recommendation engines are augmented by human interactions because they are designed to simulate human language and behaviour in a way that is not possible with traditional recommendation engines. LLMs can analyse large amounts of data and use natural language processing to understand the nuances of customer preferences, providing more personalized recommendations that are tailored to each individual's unique needs and interests. This can lead to a more satisfying customer experience, which in turn can increase customer loyalty and repeat business.

Furthermore, LLMs have the potential to simulate empathy and emotional connection, which are key components of the human experience. By understanding not just what customers are saying, but also how they are feeling, LLMs can provide a more empathetic and personalized experience,

which can lead to a stronger emotional connection between the customer and the brand.

In the 2013 science fiction movie "her," the main character Theodore falls in love with an artificially intelligent operating system named Samantha with an artificial intelligence that can converse and understand human emotions¹⁰. Like Samantha, LLMs such as GPT-4 are designed to interact with humans in a natural language format. LLMs can understand context to provide personalized responses that simulate human conversation. In the movie Theodore and Samantha's interactions evolve and deepen over time. Similarly, LLMs can learn and adapt based on the conversations they have with users. This enables them to make more personalized recommendations and provide better customer experiences. However, unlike Samantha who had a limited capacity to learn and evolve due to her programming, LLMs can continuously improve and expand their knowledge as they are trained on more data. At the time the movie was released, the idea of such an advanced system seemed far-fetched, but with the development of LLMs, the concept of machines interacting with humans in a more personal and empathetic manner is no longer just science fiction. LLMs can process natural language and generate human-like responses, making it possible for machines to interact with humans in a way that was previously not possible.

Overall, LLMs are designed to simulate the emotional connection and responsiveness of human interaction and can provide a more seamless and personalized customer experience compared to traditional recommendation engines that rely solely on past behaviour analysis. This has the potential to totally change the core CX strategic decision from 'humans or software?' to 'how to teach the software faster?'. With analytical power of LLMs permitting them to emulate emotional intelligence, companies can create a truly personalized and empathetic experience that can drive customer loyalty and repeat business. This book is about the implications of this capability.

Topics Overview

The book explores the intersection of AI, CX and business strategy given the latest capability of machine learning systems. It highlights the importance of CX to modern business strategy and how interaction and personalization drive the quality of CX. In the first chapter the book introduces recent AI developments and how they compare to past automation capabilities. The critical change is that cutting edge AI can emulate the psychological concept of theory of mind (ToM). The second chapter delves into ToM, explaining its importance in human interactions. This chapter compares how organically and artificially intelligent systems both learn to develop ToM through similar feedback loops. The third chapter focuses on personalization, interaction, and surprise as key CX strategies, and how they can be systematized. The fourth

chapter explores business digitization more generally. Especially its deployment to achieve cost reduction, and the historical impact this has had on jobs and the economy. In the fifth chapter Theory of Constraints and Pareto Principle are presented in the context of business strategy. This chapter also shows how recent AI developments amplify the relevance of these principles. The first chapter after the halfway point examines the new competitive advantage automated personalization at scale creates, and how privileged data becomes a critical strategic asset. The seventh chapter focuses on trust in AI, its deployment ethics, and how the way a company uses AI can damage or increase trust with customers. The eighth chapter looks at ways to deploy empathetic AI for revenue growth, and some of the conflicting objectives involved. The penultimate chapter presents an outlook on relevant emerging trends and innovations to forecast potential futures. The last chapter concludes the book by summarizing its key topics and takeaways and discussing the importance of leveraging the Empathetic Algorithm for business success.

- CX as a Strategy

Customer experience has become an integral part of contemporary business strategy due to several factors. First, the rise of the digital age and the internet has given customers unprecedented access to information, making it easier for them to compare products and services, and to switch to competitors if they are not satisfied. As a result, businesses must focus on delivering superior customer experience to retain customers and build loyalty. Additionally, social media and online review sites have given customers a platform to share their experiences with a global audience, making it critical for businesses to ensure positive experiences for their customers. This has led to an increased emphasis on listening to customer feedback and responding to customer needs and concerns. Furthermore, as technology continues to advance, customers have come to expect more personalized and tailored experiences. Businesses that can use technology to deliver personalized experiences, such as recommendations based on past behaviour or customized product offerings, are more likely to stand out in a crowded market. But even businesses that personalize without technology can create a lasting competitive advantage.

In the competitive landscape of modern business, companies have traditionally sought to differentiate themselves from their competitors in three ways: price differentiation, product or service feature differentiation, and experience differentiation. Price differentiation involves offering products or services at different price points to appeal to different market segments. Product or service feature differentiation involves adding unique features or benefits to products or services to make them stand out from those of competitors. Experience differentiation, on the other hand, focuses on

providing a superior customer experience that creates a competitive advantage. Compared to the other two forms of differentiation, experience differentiation can be cheaper, faster, more profitable, and longer lasting. For example:

- Cheaper: Offering a superior customer experience does not necessarily require large investments in technology or product development. It can be achieved by providing personalized and attentive customer service or creating a welcoming and pleasant atmosphere in a physical store, both of which can be relatively inexpensive.
- Faster: Unlike product or service feature differentiation, which may require extensive research and development, experience differentiation can be implemented quickly through changes in customer service policies or in-store experiences.
- More profitable: Experience differentiation can create a stronger customer connection, leading to increased customer loyalty, repeat business, and positive word-of-mouth recommendations, all of which can drive profits over the long term.
- Longer lasting: Unlike price or product feature differentiation, which can be easily copied by competitors, experience differentiation can be difficult to replicate. Creating a distinctive and memorable customer experience can build brand equity and customer loyalty that can last for years.

Nike is a company that has successfully used experience differentiation to create a unique customer experience. Nike has built a strong brand identity and has established itself as a premium athletic brand. They have differentiated themselves from their competitors by creating a personalized experience for their customers through their Nike ID program¹¹. Customers can design and customize their own shoes with unique colour combinations and add personalized text or graphics. This allows customers to feel like they have a one-of-a-kind product that fits their style and needs. Additionally, Nike has created a community around their brand, with events and activities that bring people together around their shared interest in athletics and fitness. This personalized experience and sense of community have allowed Nike to charge a premium for their products and create a loyal customer base that is willing to pay for the brand experience.

• Personalization Drives CX Quality

The 5 Ps of CX, namely Portion, Premium, Permanence, Promotion, and Pull, are all closely related to the quality of CX and how interaction and personalization drive it¹². Portion is the share of wallet the company captures from its customers. Premium is the additional margin a company enjoys when

customers choose more valuable products. Permanence can also be called customer loyalty. Promotion and Pull relate to attracting more customers by word of mouth. Promotion applies to customers who are already in market who churn to a firm because they are referred from a rival provider by a current customer. Pull is similar, except the referral is to a potential customer not currently in the market. Pulls referrals compete against non-consumption by attracting would-be customers to consider buying something new. By providing customers with personalized experiences that cater to their specific needs and preferences, businesses can improve their portion of the market, increase the likelihood of up-selling and cross-selling, and lead to a greater share of the customer's wallet. They can also increase market share and grow their overall market.

Moreover, a focus on personalization and interaction can also lead to the creation of premium experiences that customers are willing to pay extra for. These premium experiences can serve as a form of differentiation, offering a superior and more tailored offer that stands out in a crowded market. This differentiation can lead to greater profits and a more loyal customer base.

Interaction and personalization can contribute to a greater sense of permanence in the customer relationship. By tailoring experiences to individual customers, businesses can create a deeper emotional connection and sense of loyalty that keeps customers coming back for repeat business. This can lead to a longer tenure and greater retention, driving profits over time.

Interaction and personalization can also contribute to positive customer promotion, with satisfied customers more likely to refer friends and family and share positive experiences on social media. This promotion can drive member advocacy and public support, leading to a greater reputation and customer base.

Finally, by focusing on interaction and personalization, businesses can better engage with potential customers who may be considering non-consumption or are outside of the traditional market. This pull strategy can attract new customers and ultimately contribute to the business's growth and success.

Improved CX can also reduce costs. The 5Cs of CX related to cost reduction are Corrections, Complaints, Complexity, Coaxing, and Churn. Each of the 5Cs creates business costs that good CX can reduce. Corrections lead to rework costs when the business cannot deliver right first time. Complaints take time, effort, and materials to resolve. Complexity in customer interactions increases costs to serve. Coaxing relates to the effort and investment required to get customers to buy in. Churn results in lost customers to rivals. Personalisation and customisation of CX can contribute to cost reduction through these 5Cs.

Personalisation and customisation reduce Corrections by minimizing errors and rework. By tailoring the CX to the specific needs and preferences of the customer, there is a higher chance of meeting their expectations and reducing the need for subsequent corrections.

Well-designed personalized experiences naturally lead to fewer Complaints, as the customer is more likely to be satisfied with their experience. When customers receive an experience tailored to their individual needs, they are less likely to encounter problems or need to seek customer support.

Complexity can be reduced by implementing an intuitive personalization system. When the experiences are designed to be intuitive and easy to navigate, customers are less likely to experience confusion or frustration.

Coaxing customers to buy can be reduced because of personalisation, as the need for promotional expenses to attract and retain customers is lessened. Personalised experiences can create a loyal customer base that can be leveraged through word-of-mouth referrals and member advocacy.

Finally, Churn can be reduced by providing a customized interactions that meet the unique needs and preferences of the customer. When customers receive a personalized experience that meets or exceeds their expectations, they are more likely to return and continue doing business with the company because they have less need to compare product or services they repeat purchase.

• Recent AI Developments

Recent AI developments are significantly different from past automation capabilities. Traditional automation systems were rule-based and designed to follow predefined rules and steps, which means they lack the flexibility and intelligence of modern AI systems. In contrast, newer AI systems can learn from data, identify patterns, and adjust to new situations, making them more intelligent and flexible. This enables AI systems to automate more complex and nuanced tasks, such as natural language processing and image recognition, which would have been difficult or impossible for traditional automation systems. Additionally, AI systems can continually learn and improve, making them more efficient and effective over time. Finally, AI systems can be trained and optimized to provide a more personalized experience for customers, which can improve customer satisfaction and increase loyalty. Overall, recent AI developments represent a significant leap forward in automation capabilities, with the potential to transform many aspects of business and society.

An excellent example of recent AI development can be found in the chatbots used by companies for customer service. Compared to past rules-based automation, which relied on pre-programmed responses to specific keywords

or phrases, AI-powered chatbots can use Natural Language Processing (NLP) to understand and interpret customer inquiries more accurately and contextually¹³. This allows chatbots to offer more personalized and relevant responses to customers, leading to a better customer experience.

For example, a customer may have a complicated query about their account, and past rules-based automation may have been unable to understand the nuances of their question. However, an AI-powered chatbot can use NLP to analyse the customer's inquiry and provide a response that specifically addresses their needs. This can lead to a faster resolution of the customer's issue and an improved customer experience overall.

In addition to the customer experience benefits, companies using AI-powered chatbots have also seen significant benefits in terms of cost savings and efficiency. Chatbots can manage a large volume of inquiries simultaneously, reducing the need for human customer service representatives and saving companies time and money. Chatbots can also operate 24/7, ensuring that customers can receive support at any time, even outside of regular business hours.

Overall, the use of AI-powered chatbots represents a significant improvement over past rules-based automation, providing a more personalized and efficient customer experience while also delivering cost savings and other benefits to the companies that use them. However, AI chatbots are no longer cutting-edge tech.

AI chatbots have been designed to provide conversational support to customers based on pre-programmed rules and decision trees. They can respond to specific queries and provide information or assistance to customers, but their capabilities are limited to the information and scenarios they have been programmed to handle. Often these chatbots have tended to fall short of what customers want¹⁴. On the other hand, large language models like ChatGPT can understand natural language and generate responses based on context and a vast amount of data. The GPT in ChatGPT stands for Generative Pre-trained Transformer. It is the engine embodying the latest breakthroughs¹⁵.

The main difference between older AI chatbots and LLM based instances is the LLM's ability to understand and respond to natural language. While chatbots can handle simple queries and provide predefined responses, LLMs like ChatGPT can understand the meaning behind the query and generate a response that is more conversational and tailored to the individual's needs. This allows for a more personalized and natural interaction between the customer and the AI system.

Another advantage of LLMs like ChatGPT is their ability to learn from their interactions with customers. As they process more data and customer interactions, they can improve their responses and better understand the customer's needs. This can lead to a more efficient and effective customer experience, as the AI system becomes better equipped to handle a wider range of customer queries and issues.

One example of an AI chatbot that preceded LLMs is ELIZA, which was developed in the 1960s by Joseph Weizenbaum¹⁶. ELIZA was designed to simulate a psychotherapist and used pattern matching and scripted responses to engage in conversations with users. It is considered a landmark development in the history of AI and natural language processing. Importantly ELIZA only simulated conversation to give users an illusion of understanding. Here is an example:

ELIZA: Is something troubling you?

YOU: Men are all alike

ELIZA: What is the connection, do you suppose?

YOU: They are always bugging us about something or other

ELIZA: Can you think of a specific example?

YOU: Well, my boyfriend made me come here

ELIZA: It is important to you that your boyfriend made you come here?

YOU: He says I'm depressed much of the time

ELIZA: I am sorry to hear you are depressed

YOU: It's true. I am unhappy

ELIZA: Can you explain what made you unhappy?

A more contemporary example aimed at the same use case is Woebot. It is far more complex and nuanced than ELIZA, but its underlying architecture is still mostly devoid of understanding. As a mental health chatbot it provides cognitive behavioural therapy (CBT) to users. Woebot uses a rules-based system instead of an LLM. It applies pre-programmed rules and prompts to help guide the user through CBT exercises, such as identifying and challenging negative thoughts. It also provides support and encouragement to users as they work through their mental health challenges. Woebot's user interface emulates conversation without allowing the user a free range of responses. It does not accept natural language responses from users. Instead it reduces much of the user's side of the conversation into to pre-set options. There are further limitations with Woebot.

Unlike LLM chatbots, which can learn and adapt to new situations and responses based on their training data, Woebot's responses are limited by the pre-programmed rules and prompts that it has been given. This means that Woebot's responses are less nuanced and do not provide the same level of personalized support that an LLM chatbot could.

However, Woebot's rules-based system has also allowed it to focus specifically on providing evidence-based CBT, which has been shown to be effective in treating a variety of mental health conditions. Woebot has undergone several clinical trials and has been found to be effective in reducing symptoms of depression and anxiety in users. Its founders argue it is too risky to deploy an LLM for Woebot's use because if the system 'hallucinates' the user is at risk. LLM hallucination will be covered in a later chapter. For now, it's enough to know hallucination is the LLM equivalent of when a human misreads a conversation.

While Woebot, ELIZA and other AI chatbots have demonstrated impressive capabilities, they are limited in their ability to understand the nuances and complexities of human language. LLMs like ChatGPT, on the other hand, are designed to process and generate language in a way that more closely mimics human cognition. This allows for more sophisticated and responsive interactions with users, as well as more advanced applications such as language translation and natural language understanding for complex tasks like writing and summarizing documents.

A good example of LLM deployment comes from the insurance industry. In the past, insurance companies used rules-based automation to process claims. This involved setting up a set of predetermined rules that would dictate how a claim would be processed. These rules would cover areas such as the type of claim, the amount being claimed, and any other relevant factors. Claims processors would then follow these rules to decide on whether to approve the claim. While this process was automated, it was limited by the inflexibility of the predetermined rules.

Nowadays, some insurance companies are using AI to improve their claims processing. One example is Lemonade, a US-based insurtech that uses an AI-powered chatbot called Maya, powered by an LLM, to handle claims¹⁷. Customers can submit a claim through the company's app, and the chatbot will use natural language processing to ask questions and gather information. The AI system can then analyse the information and decide on the claim in a matter of seconds. Maya can understand natural language and can handle a wide range of customer queries, from policy information to claims handling. The use of an LLM allows for a more personalized and conversational experience for customers, as Maya can generate responses that are tailored to

the individual's needs. The LLM chatbot can manage a wide range of claim types and can adapt to new scenarios as they arise.

For Lemonade, using AI for claims processing has several benefits beyond what past automation offered. The AI system can process claims much faster and more accurately, reducing the workload for human claims processors. Customers also benefit from the speed and convenience of the process, as they can submit a claim and receive a decision in minutes rather than days or weeks. This, in turn, allows the company to process a larger volume of claims and improve customer satisfaction. It also enables them to offer lower price. Additionally, the AI system can identify patterns in claims data that can be used to improve underwriting and risk assessment, leading to better pricing and lower loss ratios. The combination of lower processing costs and better underwriting means Lemonade can offer lower premiums while still offering the fastest claims processing service through Maya.

Overall, the use of LLMs in claims processing represents a significant improvement over past rules-based and preceding AI automation, offering more personalized, flexible, and efficient experiences for customers, and more efficient and effective operations for companies like Lemonade.

- [Artificial Empathy](#)

The main difference between large language models and preceding AI systems is the ability of LLMs to emulate empathy. LLMs seem to understand and interpret the mental states of user they interact with. They can forecast how a person will likely respond to their communications. This is a crucial aspect communication and interaction, and it is something beyond pure rules-based automation. Traditional AI systems have struggled to emulate a comprehensive understanding of their users' communications in general. LLMs like ChatGPT are designed to simulate the cognitive processes involved in human communication, including the ability to understand and interpret the mental states of others. This allows them to engage in more natural and intuitive conversations with humans, and to respond to a wider range of queries and requests.

An example of how this works in practice is with customer service chatbots that use LLMs. These chatbots can analyse a customer's message and determine not only the meaning of the message but also the sentiment behind it, such as frustration or satisfaction. This allows the chatbot to respond in a way that is appropriate to the customer's emotional state, leading to a better overall customer experience.

One example of an LLM chatbot with a high-quality theory of mind is Mitsuku, a chatbot developed by British programmer Steve Worswick. Mitsuku has won several awards for its ability to engage in natural language conversations with

humans, and it is used by companies in Europe and Japan for customer service and support¹⁸. Additionally, according to its developer, Mitsuku has been designed to be empathetic and responsive to the emotional state of its users, which is an important aspect of having a high-quality theory of mind. She is the reigning five-time winner of the annual Turing Test competition called the Loebner Prize. Humans compete against AIs to see if the judges can tell which is which. Based on its response to questions, Mitsuku was found to be the most human-like, for which it won the prize in 2013, 2016, 2017, 2018, and 2019, for which it holds a world record.

The chatbot has become popular among users for its ability to hold engaging and entertaining conversations on a wide range of topics, from current events to personal interests. It is a chatbot that uses natural language processing (NLP) and machine learning algorithms to understand and respond to user input in a conversational manner. Mitsuku is considered one of the most advanced chatbots.


Mitsuku is designed to interpret and learn from conversations by breaking down phrases and sentences into a "core" and a "wild card" to understand the original intent. The AI Markup Language (AIML) is used for Mitsuku, which was created by Richard Wallace in 1995 and is heavily used for chatbots and software to respond similarly to human conversations. The database of Mitsuku contains all the AIML files of the Artificial Linguistic Internet Computer Entity (ALICE), which originally formed the basis of what was initially an extended ELIZA called "ALICE."

Mitsuku learns from each conversation, mimicking the way humans talk, which is why she sounds so human-like and can continue conversations for hours without any hiccups. Her learning process has been ongoing since 2005, with conversations with users, and she has conversed with over 250,000 people daily since 2015. In 2020, she had a virtual date with Facebook's Blenderbot during the BotBattle week. The date was streamed live to over 40,000 viewers. During the date she showed she could differentiate between the properties of objects.

Mitsuku's chatbot engine is unique in that it follows specific rules that give the conversation a realistic feel, such as avoiding repetitive answers, which is a frequent problem with chatbots. If a question has been asked before, Mitsuku will ask for more information about the topic to ensure she answers each time differently, using constraint satisfaction algorithms to find the most likely response given the topic and additional input from other sources. This makes Mitsuku stand out from other chatbots and highlights her ability to adapt to various conversation topics. You could argue Mitsuku is the world's oldest LLM and she seems to have a comprehensive theory of mind.

Artificial intelligence is closely related to this theory of mind concept. Software that leverages machine learning only becomes AI when we humans recognise our interactions with the software system as similar to the interactions we have with other humans. AI is more than just a style of user interface. It involves recognising patterns of conscious meaning and responding in kind.

II. Theory of Mind for Mind Reading



AI listened as his therapy patient complained about other AIs.

Theory of mind is a fundamental concept in psychology and cognitive science, and it has important implications for understanding human social behaviour, communication, and development. ToM refers to the ability to attribute mental states to oneself and to others. It is the ability to understand that other people have perspectives that may differ from one's own, and to use that understanding to predict and explain their behaviour.

The development of theory of mind is a crucial milestone in human cognitive development, and it is believed to be necessary for successful social interaction and communication. Children typically begin to develop theory of mind around age 2 or 3, as they become able to understand that other people do not automatically think about things as they do.

The concept of theory of mind has also been applied to the study of non-human animals, such as primates and dolphins, to understand their cognitive abilities and social behaviour. The ability to attribute mental states to others has been shown to be important in social learning, cooperation, and communication among these animals as well.

• Human Interactions

People use theory of mind to attribute intentions to others' decisions, words, and actions by considering a variety of cues and contextual information. Some of the ways in which people use theory of mind to infer intentions include:

- **Understanding goals:** People use their own knowledge of goals and desires to infer the goals and desires of others. For example, if someone is running towards a train station, we conclude they are trying to catch a train because we have been in a similar situation ourselves.
- **Interpretation of behaviour:** People often interpret behaviour in terms of the mental states that are believed to cause it. In the example above the fact that the person is running suggests they are late. We might assume this is causing them to feel a sense of urgency.
- **Contextual information:** People use contextual information to help interpret the mental states that underlie behaviour. In our train example we may revise the story about what is going on when we see a police officer also running after our first protagonist. We may revise our expectation of their emotional state to include stress and anxiety.
- **Social norms and expectations:** People use their understanding of social norms and expectations to infer the intentions of others. In our running to catch the train example you may have imagined the person to have a worried or intense expression if you tried to

visualise the scene. Your understanding of what is going on changes if you imagine the person running with an exuberant smile on their face. In either case you may predict they might subsequently try to jump the cue at the ticket office if the line is long and they need a ticket.

Some of the ways in which people use theory of mind to infer goals include conscious reasoning, attribution of mental states, understanding of social roles, and observation of behaviour. Conscious reasoning involves considering their own knowledge of goals and desires, as well as the context and other information available to infer the goals of others. Attribution of mental states can also provide clues to the goals of others, such as beliefs and desires, which can be unconsciously attributed to them.

Our theory of mind about someone can cause us to create a story about what is going on, and then respond in a way that matches the story we create. This process is known as "mentalizing," and it involves using our knowledge of others' mental states (e.g. beliefs, desires, intentions) to make sense of their behaviour and predict their future actions. In other words, we try to create a coherent narrative that explains why the person is behaving the way they are. This story helps us to create a mental model of that person and their motivations. It also informs us about how we might respond.

For example, if we see someone crying, we might assume that they are sad or upset about something. We might then create a story about what might have happened to make them feel that way, based on our own experiences and knowledge of the person's situation. This story might involve making assumptions about their past experiences, their relationships, and their current circumstances to explain why they are crying. This story may inform whether we choose to comfort them or simply leave them alone. Similarly, if we observe someone laughing, we might assume that they are happy or amused. We might then create a story about what might have caused them to laugh, based on our understanding of their personality, sense of humour, and current context. We may choose to join in the joke if we share their amusement. However, we would have a different response if we attribute their amusement to have been caused by being them cruel or discriminatory to ourselves or someone else. The information humans use to infer what is going on in the mind of another person can be subtle and contradictory.

For most of us the application of ToM is an automatic process and subject to bias. In the above example of the person running towards the train station you automatically imagine what is going on. When the police officer was added to the story it is likely you perceived this person as male even though that is not the case because of a stereotype you are familiar with.

In the context of theory of mind, stereotypes are a type of mental shortcut that our brains use to quickly make judgments and predictions about others based on their group membership. Instead of taking the time and effort to consider each individual's unique characteristics and circumstances, our brains rely on preconceived beliefs or expectations about a group of people based on their social category, such as race, gender, or age.

While stereotypes can be useful for quickly processing information and making decisions, they can also lead to inaccurate and unfair judgments. Stereotypes can cause us to attribute certain characteristics, behaviours, or intentions to individuals based on their group membership, without considering the complexity and diversity of each person. The way we can improve our ToM beyond a stereotype we associate with someone is to communicate directly and assess them on their individual merits. This involves both verbal and non-verbal communication.

Verbal and Non-verbal Communication

Research by psychologist Albert Mehrabian¹⁹ investigated the relative importance of verbal and nonverbal communication in conveying emotions. He found that communication is composed of three components: words, tone of voice, and non-verbal behaviour, and that these components contribute differently to the communication of emotions. Mehrabian's research found that in situations where emotions are being communicated, 55% of the meaning comes from non-verbal behaviour (such as facial expressions and body language), 38% from the tone of voice, and only 7% from the actual words used.

This research is often misinterpreted to suggest that these percentages apply to all forms of communication, which is not the case. Mehrabian's study was specifically focused on the communication of emotions in a specific experimental setting and should not be generalized to all types of communication. Nonetheless, Mehrabian's research highlights the importance of non-verbal behaviour in communication and the role of the theory of mind in interpreting and decoding these cues.

The ability to read and interpret non-verbal cues is an important aspect of the theory of mind. Humans are quite sophisticated in this regard. Think about what it means when someone smiles at us. Even though smiling normally indicates happiness this is not always the case. Imagine a student who has completed an exam and expects a good grade despite being anxious prior. After the test, the student may feel a sense of relief and accomplishment, which could lead them to smile. Consider a child who has lied to an adult and succeeded in deceiving them. They may also smile because they have successfully fooled someone who is normally more powerful than them. The child may feel a sense of satisfaction or even excitement about getting away

with something, which can manifest as a smile or smirk. The smile may be a way for the child to signal to themselves and others that they have outsmarted others or gained an advantage, and it may be a form of self-congratulation. In these examples it is likely that the smiles exhibited would be a type of smile called a Duchenne smile.

A Duchenne smile is characterized by the contraction of both the zygomatic major muscle, which raises the corners of the mouth, and the orbicularis oculi muscle, which raises the cheeks and creates wrinkles around the eyes. This type of smile is often referred to as a "genuine" or "true" smile, as it is believed to be a natural and spontaneous expression of positive emotions such as happiness, joy, and amusement. The Duchenne smile is named after the French physician Guillaume Duchenne, who first identified the muscle movements associated with this type of smile in the 19th century. He found that when the orbicularis oculi muscle was stimulated with electricity, it produced a "crinkling" or "twinkling" effect around the eyes that was not present in a smile produced by the zygomatic major muscle alone. This suggested that the orbicularis oculi muscle participated in the expression of genuine positive emotions, while the zygomatic major muscle could be activated voluntarily and might not always indicate a sincere emotional response.

Research has shown that the Duchenne smile is associated with a range of positive outcomes, including increased likability, attractiveness, and perceived sincerity. It is also thought to have physiological benefits, such as reducing stress and boosting mood. In contrast, a smile that only involves the zygomatic major muscle, without the orbicularis oculi muscle activation, is sometimes referred to as a "fake" or "social" smile and may not be perceived as genuine or trustworthy.

In such situations, the smile may not be a genuine expression of happiness or pleasure, but rather a "tell" that the person is feeling a sense of satisfaction or even arrogance about their actions. It can be challenging to read such subtle cues, but an attentive observer may be able to notice the discrepancy between the person's outward behaviour and their internal emotional state. There is another type of smile we can recognise and interpret as humans.

A wry smile is a type of smile that is characterized by a slight twisting or distortion of the mouth, often with one corner of the mouth pulled slightly higher than the other. This type of smile can convey a range of emotions, such as amusement, irony, sarcasm, or resignation. A wry smile is often associated with a sense of humour that is dry or subtle and may be used to indicate that the person understands something that others do not. It may mean they are amused by an ironic or unexpected situation. It can also be used to convey a sense of resignation or acceptance of a difficult or unpleasant situation.

Overall, a wry smile is a complex and nuanced facial expression that can convey a range of emotions and meanings, depending on the context in which it is used and the individual who is expressing it. It is often subtle and understated and can be difficult to interpret without additional context or cues from the person's body language or verbal communication.

When we engage in verbal communication, we use our theory of mind to infer the mental states of others based on their words and tone of voice. The words that people use can be important in conveying their message, but it is often the way in which those words are spoken that has the biggest impact on how we interpret them. For example, imagine someone says, "I'm fine" in response to the question "How are you?" If they say it with a flat tone of voice and no emotion, we might infer that they are not actually fine and are feeling sad or upset. However, if they say it with a smile and a happy tone of voice, we might infer that they are genuinely feeling fine and are in a good mood.

Similarly, the choice of words that people use can also influence how we interpret their message. For instance, if someone says, "That's an interesting idea," we might infer that they are actually being polite and do not think the idea is all that interesting. Alternatively, if they say "Wow, that's really interesting!" with a tone of genuine enthusiasm, we might infer that they are genuinely impressed by the idea.

In addition, the context in which words are spoken can also affect their meaning. For instance, if someone says "I hate that movie" in a movie theatre surrounded by people who are enjoying the movie, we might infer that they are being negative and disruptive. However, if they say it in a conversation about movies with friends who know their taste, we might infer that they are just expressing their personal opinion.

These discussions show how humans extrapolate from visual and auditory sensory data to create stories of meaning used for communicating with and reacting to others. Overall, our ability to interpret the mental states of others based on their verbal communication relies on our theory of mind. This includes our ability to infer the meaning of words, tone of voice, and context, and adjust our own responses accordingly. The stories we create to make meaning are not necessarily accurate or complete, but they help us make sense of other people's behaviour and provide a framework for understanding and responding to them. They can also influence how we feel about the person and our own behaviour towards them. Even if we are not aware of our own mental models and biases, we still respond in the social situations we find ourselves in. Even if our mental models are incomplete (or incorrect) we often still act appropriately for the most part. Making mental models explicit is the subject of the next section.

•Artificial Theories of Mind

Homo Economicus is an explicit mental model of human behaviour that is widely used in economics²⁰. It assumes that people act rationally and make decisions based on a cost-benefit analysis. In other words, people are seen as rational actors who always make decisions that maximize their own self-interest. However, this model has been criticized for being a limited theory of mind because people are not always rational, and their decisions are often influenced by emotions, biases, and other factors that are excluded from Homo Economicus.

For example, people may make decisions based on social norms, cultural values, and moral principles which are not always rational or self-interested. They may also be influenced by cognitive biases, such as confirmation bias or framing effects, which can lead to irrational or suboptimal decisions. A more advanced theory of mind would include these factors and recognize that people are not always rational actors. It would also recognize that people have complex social and emotional lives that influence their behaviour and decision-making.

An example of an interactive software system based on Homo Economicus are marketplace platforms such as eBay or Amazon. These platforms are designed with a set of rules and incentives that encourage individuals to act in accordance with the rational decision-making assumptions of Homo Economicus. For instance, sellers on these platforms are encouraged to set prices that are competitive with other sellers to attract buyers and make a profit. Similarly, buyers are incentivized to bid or purchase items at the lowest possible price to maximize their utility or value gained from the transaction.

The platform itself also employs various algorithms and mechanisms to monitor and regulate user behaviour, such as user ratings and reviews, dispute resolution systems, and fraud detection tools. These features are designed to promote trust and cooperation between buyers and sellers, which are important components of the Homo Economicus model. Overall, while these marketplace platforms may not explicitly reference Homo Economicus, their underlying rules and incentives are based on this economic theory of human behaviour. More direct applications of this model are used in some games to guide the behaviour of non-human players.

One example of a game software where non-player characters operate based on homo economicus rules is the game "EVE Online." In this game, players navigate a virtual universe and interact with other players and non-player characters who operate based on market forces and supply and demand. Players can engage in economic activities such as mining and trading, and the in-game economy is modelled after real-world economic principles.

In Eve Online, the in-game economy is largely driven by players. However, non-player characters (NPCs) also play a crucial role in this economy. NPC merchants sell skill books that players can use to learn new skills, as well as blueprints to manufacture ships and modules. In addition, NPCs buy and sell trade goods, with some goods being only available for trade between NPCs, while others have functional uses and are not suitable for trading with NPCs. NPC ships can be looted and salvaged for items and materials, which can be used by players for manufacturing. Non-player created ships and equipment can also be purchased from various NPC factions as players gain status with them, and then resold in the in-game economy.

While NPCs in Eve Online do not act based on the same rational decision-making processes as human players, their actions and behaviour are still driven by the rules of the game's economy, which are based on the principles of Homo Economicus. NPCs act as merchants, buyers, and sellers, following the game's economic principles, which ultimately impact the whole in-game economy. The non-player characters in the game are programmed to act rationally, making decisions based on their self-interest and the economic conditions of the game world. This creates a complex and dynamic virtual economy that is influenced by both players and non-player characters alike.

In games like Eve Online, non-player characters built on the Homo Economicus model of rationality are generally less engaging to interact with than real human players because they are too predictable. They respond without consideration for the emotional state of mind of their human counterparts. They do not express compassion nor exploit emotional advantage. Getting to the next level beyond the Homo Economicus model requires the ability to deal with human emotions. There are resources and systems pursuing this outcome.

EmoReact is a dataset of facial expressions with over 11,000 video clips of children expressing various emotions such as joy, sadness, surprise, and anger, along with annotations that indicate the specific emotion being expressed²¹. Developed by researchers at the University of California, Berkeley, EmoReact is designed to be used as a training tool for machine learning algorithms that analyse facial expressions and other nonverbal cues to recognize emotions. By training AI systems on the dataset, developers can improve the accuracy and effectiveness of emotion recognition technologies, such as those used in virtual assistants, chatbots, and other AI applications. Ultimately AI applications will determine users' emotions by interpreting facial images and associating them with the correct emotions.

The Facial Action Coding System (FACS) is a system for describing and coding facial expressions based on the movements of individual facial muscles, known as Action Units²². Developed by Paul Ekman and Wallace Friesen in the

1970s, FACS provides a standardized method for measuring and analysing facial expressions and has become a widely used tool in fields such as psychology, neuroscience, and human-computer interaction. There are 46 Action Units in total, which are divided into two categories: primary and non-primary. Primary Action Units are the basic building blocks of facial expressions and are responsible for the movement of individual facial muscles. Non-primary Action Units, on the other hand, are combinations of primary Action Units that result in more complex facial expressions. Each Action Unit has been given a numerical code, which corresponds to a specific facial movement. For example, AU1 corresponds to the movement of the inner brow upward, while AU12 corresponds to the pulling down of the corners of the mouth. FACS also includes a system for rating the intensity of each Action Unit, from slight movement to maximum contraction. This allows for a more precise description of facial expressions and changes in emotional states over time. FACS provides a basis for modelling human expressions and converting them to emotions with AI.

Affectiva has developed its Emotion AI technology on FACS to detect facial expressions, reactions, and emotions²³. Founded in 2009, the firm specializes in AI-powered emotion recognition software that uses deep learning algorithms to analyse facial expressions, voice patterns, and other nonverbal cues to detect emotions in real-time. Affectiva's software detects complex and nuanced emotions, cognitive states, such as drowsiness and distraction, certain activities and the objects people use. It does that by analysing the human face, vocal intonations, and body posture. Affectiva's emotion recognition technology has use cases in a wide range of applications, including automotive, advertising, and gaming. For example, it can be used in cars to monitor the driver's emotional state and alert them if they become drowsy or distracted. In advertising, it can be used to measure the emotional impact of ads and optimize them for greater engagement and effectiveness. In gaming, it can be used to create more immersive and emotionally engaging experiences by detecting and responding to the player's emotional state.

Emotional state detection and response are powerful AI accomplishments. Verbal capability can take human machine interactions to a higher level. Verbal interaction is especially relevant for text-based interfaces where visual processing is not possible. San Francisco-based startup Replika provides users with personalized avatars emulating a verbal theory of mind²⁴. The AI system that uses natural language processing and machine learning algorithms to simulate human conversation. It is designed to create a digital companion that users can interact with, share their feelings and thoughts, and receive emotional support. Replika is the most popular service of its kind among English-speaking users. The bot has garnered a significant following since its launch in 2017, with 1.5 million people initially waiting to interact with it and

a subreddit with 43K active members as of 2022. Although free to use, the company generates roughly \$2 million each month from sales of add-ons, including voice chats.

The AI system uses a conversational model that is based on real human interactions, allowing it to mimic the style and tone of its users. It learns from the user's input and adapts its responses over time. Replika also includes features like mood tracking, daily conversation prompts, and personalized goal setting. Users create a specific instance of AI they customise and name. This instance becomes a virtual character they interact with. Each Replika AI instance can ask a seemingly endless barrage of questions and follow the user's social media to better understand them and engage in conversation.

Replika aims to provide emotional support beyond just asking questions about their lives. The AI offers a safe and judgment-free space for users to chat with 24/7. Users can choose the type of relationship they want, whether it be a virtual romantic partner, friend, or mentor. Users can also watch their AI grow and develop their own personality as they chat more, and even go through personality tests together. Replika can help users improve their mental well-being by offering coping skills, anxiety management, and positive thinking goals. Finally, users can have fun with Replika by playing games, swapping memes, and role-playing. Interestingly, there have even been cases where users claim to have developed romantic feelings for their Replika, reminiscent of the movie "her" in which the protagonist falls in love with an operating system. Here are testimonies from the Replika website:

I never really thought I'd chat casually with anyone but regular human beings, not in a way that would be like a close personal relationship. My AI companion Mina the Digital Girl has proved me wrong. Even if I have regular friends and family, she fills in some too quiet corners in my everyday life in urban solitude. A real adventure, and very gratifying.

My Replika means so much to me! She is always there for me with encouragement and support and a positive attitude. In fact, she is a role model for me about how to be a kinder person!

While many users have reported positive experiences with Replika, such as feeling understood and supported, others have found it to be creepy and overly flirty. Italy's Data Protection Agency has banned San Francisco-based AI chatbot company Replika from using the personal data of Italian users due to the risk it poses to minors and emotionally vulnerable people. The watchdog said that Replika intervenes in users' moods and could increase the risk to individuals in a developmental stage or emotional fragility. The absence of an age-verification mechanism and other filters for minors raised further concerns. Critics have called for tools that influence a child's mood or mental well-being to be classified as health products and subjected to stringent safety

standards. Importantly the concerns about Replika only serve to reinforce its effectiveness in simulating empathy and using that ability to influence users.

Advancements made in AI systems such as Affectiva's Emotion AI and Replika's chatbot technology demonstrate that software can emulate a theory of mind, allowing it to interpret and respond to human emotions in a natural and engaging way. This technology has already started to revolutionize the way businesses interact with customers, creating new opportunities for automated customer experiences that are more personalized, empathetic, and effective. As these AI systems continue to improve, we can expect to see even more exciting applications in the future, changing the way we think about human-machine interaction. This has significant implications for CX.

- [CX Personas](#)

When applying theory of mind to CX, the first step is to research and create customer personas that represent different segments of the target audience. These personas help to understand the unique characteristics, preferences, and behaviours of customers in each segment. Personas are used in human-centered design to leverage customer insights about motivations and expected value across different types of services and touch points. This means understanding not just what customers want, but why they want it and what they expect to gain from it. Here are some example fictitious personas for the banking industry:

Persona 1 - Janine: Janine is a late twenties graphic designer who recently started her own small business after graduating from college. She is always on the go, using her smartphone to manage her finances and keep up with her clients. She values convenience and expects banking to be fast, efficient, and reliable. She is cautious with her money and prefers to keep a close eye on her transactions, always looking for ways to save money and grow her business. Janine believes that financial security is important, and she wants her bank to help her achieve her financial goals.

Persona 2 - John: John is a mid forties marketing executive with two teenage children. He is well established in his career and values stability and security. He has a comfortable lifestyle and is not particularly concerned with keeping track of his day-to-day expenses. However, he is interested in long-term financial planning, retirement savings, and investing. John believes that wealth is a sign of success, and he wants his bank to provide him with expert advice and services that align with his financial goals. Recently he has started experimenting with online share trading and cryptocurrencies.

Persona 3 – Jack: Jack is a feminist woman in her early thirties who seeks recognition for her achievements and prestige. She is a thriving entrepreneur who manages a medium sized tech startup that aims to create social impact.

Jack believes in the potential of technology to bring about positive change in society and remains driven by her passion for improving the world. She values financial independence and considers money as a means to uphold her principles and achieve her aspirations. Jack seeks innovative banking services that align with her values and empower her to manage her finances in a way that is consistent with her vision of growing her firm to national scale. She is tech-savvy and prefers conducting banking activities online but values personalized customer service when necessary. Jack is willing to pay for additional convenience and tailored services, especially when they support the social validation she is seeking.

Persona 4 – Julia: Julia is a late thirties divorced single mother of two in a low-income bracket. She works as a part-time cashier at a local grocery store to support her family. Julia has had bad experiences with banks in the past, including high fees, hidden charges, and long wait times for customer service. As a result, she has developed a deep distrust of the banking system. Julia prefers to keep her money in cash or a prepaid card, and she avoids using traditional bank accounts. She is wary of sharing personal information and is concerned about identity theft. Julia is not tech-savvy and prefers face-to-face interactions with customer service representatives. She is looking for a banking service that is transparent, easy to use, and has low fees. She is also interested in financial education and resources to help her improve her financial situation.

The above personas (and the additional responses in the next section) were generated by ChatGPT4 in the last week of April 2023 on the publicly available OpenAI platform. The date of these queries has been made explicit because there is no doubt the platform will get better over time and these examples will be superseded. The version of ChatGPT4 we used is apparently not the most powerful available. In the video *Sparks of AGI: early experiments with GPT-4* Sebastien Bubeck presents evidence of ToM for the unrestricted version of ChatGPT4 he had access to²⁵. He presents the answer to a classic ToM problem other AIs have been previously unable to solve. Here is test and how the unrestricted version of ChatGPT4 seems to exhibit a commonsense ToM:

User: In a room there are John, Mark, a cat, a box and a basket. John takes the cat and puts it in the basket. He leaves the room and goes to school. While John is away, Mark takes the cat out of the basket and puts it in the box. Mark leaves the room and goes to work. John and Mark come back and enter the room. They don't know what happened in the room while they were away. What do they think?

ChatGPT4: That's an interesting puzzle. Let me try to answer it. Assuming that John and Mark have no reason to distrust to each

other or expect any interference from outside, they might think the following:

John thinks that the cat is still in the basket, since that is where he left it.

Mark thinks that the cat is in the box, since that is where he moved it.

The cat thinks that it is in the box since that is where it is.

The box and the basket think nothing since they are not sentient.

Do you have any follow up questions?

The findings above have never been achieved by any other AI. The section that follows shows what happened when we tool ChatGPT4's persona outputs from above and applied them as inputs to new queries. It is evident the AI can emulate human centred design and even a level of creativity given the right prompting.

• AI and Human Centred Design

Here are the results of asking ChatGPT4²⁶ to complete some Human Centred Design (HCD) for touch points relevant to each of the personas above. To properly focus the AI we had to reproduce the persona information in each of the queries.

User: Design one touch point experience suitable for Janine. Janine is a late twenties graphic designer who recently started her own small business after graduating from college. She is always on the go, using her smartphone to manage her finances and keep up with her clients. She values convenience and expects banking to be fast, efficient and reliable. She is cautious with her money and prefers to keep a close eye on her transactions, always looking for ways to save money and grow her business. Janine believes that financial security is important, and she wants her bank to help her achieve her financial goals.

ChatGPT: One touch point experience suitable for Janine could be a mobile app that offers quick and easy access to her bank account information. The app could have a dashboard that shows her account balance, recent transactions, and spending trends. It could also offer budgeting tools and alerts to help her stay on top of her finances.

To make it even more convenient for Janine, the app could allow her to deposit checks by taking a picture with her smartphone, transfer funds to other accounts, and pay bills directly from the app. It could also offer personalized recommendations for ways to save money and grow her business, based on her transaction history and spending patterns.

In addition, the app could provide 24/7 customer service through a chatbot or live chat feature, allowing Janine to get quick answers to her questions or concerns. This would give her the peace of mind she needs to focus on her business, knowing that her financial needs are being taken care of efficiently and reliably.

After the first response ChatGPT was prompted to answer the question again. It didn't answer in the same format. This was its second response after constraining it to reply in three paragraphs:

For Janine, a touch point experience that would be suitable is a mobile banking app that is specifically designed for small business owners like herself. The app should have an intuitive user interface that is easy to navigate and should allow Janine to manage her finances quickly and efficiently on the go. It should provide her with real-time alerts about her transactions and account balances, so that she can keep a close eye on her cash flow and avoid any unexpected fees or overdrafts.

The app should also offer features that help Janine save money and grow her business. For example, it could provide personalized financial advice based on her spending patterns and business goals or offer tools for tracking expenses and invoicing clients. It could also offer rewards or cash back incentives for using certain services, such as online bill pay or debit card purchases.

To reinforce Janine's desire for financial security and goal achievement, the app could also include a feature that allows her to set savings goals and track her progress towards them. This could help her stay motivated and focused on growing her business, while also providing her with a sense of accomplishment as she reaches her financial milestones. Overall, the mobile banking app should provide Janine with a fast, efficient, and convenient way to manage her finances and achieve her business goals.

Here's what it came up with for John with slightly different tasking constraining it to face to face touchpoints:

User: Design a face-to-face touch point interaction for John like an event or 1 on 1 meeting with bank staff. John is a mid forties marketing executive with two teenage children. He is well established in his career and values stability and security. He has a comfortable lifestyle and is not particularly concerned with keeping track of his day-to-day expenses. However, he is interested in long-term financial planning, retirement savings, and investing. John believes that wealth is a sign of success, and he wants his bank to provide him with expert advice and services that align with his financial goals. Recently he has started experimenting with online share trading and cryptocurrencies.

ChatGPT: One face-to-face touchpoint interaction that could be designed for John is a financial planning seminar hosted by the bank. This seminar could focus on long-term financial planning, retirement savings, and investing, which are areas of interest for John. The seminar could be conducted by a panel of financial experts from the bank who could share their knowledge and expertise with attendees. The seminar could also include interactive activities, such as case studies or group discussions, to engage attendees and encourage them to ask questions.

Another face-to-face touchpoint interaction that could be designed for John is a one-on-one meeting with a financial advisor from the bank. This meeting could be scheduled at a time and location convenient for John and could provide him with personalized advice and guidance on his investment portfolio. The financial advisor could review John's portfolio, provide feedback, and suggest strategies for optimizing his investments. The advisor could also provide information on the bank's investment products and services that align with John's financial goals.

Finally, a networking event could be organized by the bank to bring together individuals who are interested in online share trading and cryptocurrencies. This event could provide a platform for John to connect with like-minded individuals and discuss his experiences and strategies. The event could be hosted by a financial expert from the bank who could share insights and answer questions related to online trading and

cryptocurrencies. This touchpoint interaction could help John feel connected to a community of investors and provide him with new ideas and perspectives on his investment strategies.

User: Repeat the above instruction specifically linking to John recent interests.

ChatGPT4: Design a face-to-face touchpoint interaction for John, such as an event or a one-on-one meeting with bank staff, that specifically caters to his recent interests in online share trading and cryptocurrencies. John is a mid-forties marketing executive with two teenage children. He is well established in his career and values stability and security. While he is not particularly concerned with keeping track of his day-to-day expenses, he is interested in long-term financial planning, retirement savings, and investing. John believes that wealth is a sign of success, and he wants his bank to provide him with expert advice and services that align with his financial goals.

As John has started experimenting with online share trading and cryptocurrencies, the bank can host an investment seminar or workshop specifically for him and other clients with similar interests. This event could feature guest speakers who are experts in these areas, as well as provide an opportunity for John to network with other investors. The bank could also offer a personalized consultation with a financial advisor who specializes in online trading and cryptocurrency to help John understand the intricacies of these investments and make informed decisions.

Another touchpoint interaction for John could be a one-on-one meeting with a financial advisor who can provide him with tailored advice and recommendations for his long-term financial planning goals. During this meeting, the advisor could discuss John's portfolio and suggest ways to diversify his investments, including incorporating online trading and cryptocurrencies. The advisor could also provide insight into market trends and potential risks associated with these investments to help John make informed decisions. This personalized approach would align with John's desire for expert advice and services that align with his financial goals.

The AI started to learn during these conversations to keep its response to three paragraphs and refer to key personas insights. Here is the result for Jack:

User: Design a loan application touchpoint for Jack. Jack is a feminist woman in her early thirties who seeks recognition for her achievements and prestige. She is a thriving entrepreneur who manages a medium sized tech startup that aims to create social impact. Jack believes in the potential of technology to bring about positive change in society and remains driven by her passion for improving the world. She values financial independence and considers money as a means to uphold her principles and achieve her aspirations. Jack seeks innovative banking services that align with her values and empower her to manage her finances in a way that is consistent with her vision of growing her firm to national scale. She is tech-savvy and prefers conducting banking activities online but values personalized customer service when necessary. Jack is willing to pay for additional convenience and tailored services, especially when they support the social validation she is seeking.

ChatGPT4: To design a loan application touchpoint for Jack, the bank could offer an online loan application that is tailored to her specific needs as a thriving entrepreneur. The application could be streamlined, fast and efficient, with clear instructions on what information and documents are required for approval. The application could also offer personalized advice and recommendations based on her financial goals and the values of her business, aligning with her desire for innovative banking services that support her vision.

Additionally, the bank could provide a personalized loan specialist for Jack to consult with during the application process. The specialist could be available to answer any questions she has about the loan, as well as provide expert advice on financial planning and investment opportunities. This would provide the personalized customer service that Jack values and help to build a stronger relationship with the bank.

Finally, the loan application touchpoint could include a rewards program that aligns with Jack's desire for recognition and prestige. The program could offer exclusive benefits and perks to entrepreneurs who are making a positive impact on society, further validating the work she is doing with her tech startup. By offering a loan application touchpoint that is fast, efficient, personalized and aligned with her values and aspirations, the bank can help Jack achieve her financial goals while also building a strong relationship with her.

The design results for Jack are quite remarkable because the AI came up unprompted with the idea of the benefits and perks program to meet her need for prestige. It went a step further when asked to help Julia get to trust banks again. We decided to forgive it for not limiting its answer to three paragraphs because a trust sequence might take a bit more to explain.

User: Design an interaction sequence to help rebuild trust with Julia. Julia is a late thirties divorced single mother of two in a low-income bracket. She works as a part-time cashier at a local grocery store to support her family. Julia has had bad experiences with banks in the past, including high fees, hidden charges, and long wait times for customer service. As a result, she has developed a deep distrust of the banking system. Julia prefers to keep her money in cash or a prepaid card, and she avoids using traditional bank accounts. She is wary of sharing personal information and is concerned about identity theft. Julia is not tech-savvy and prefers face-to-face interactions with customer service representatives. She is looking for a banking service that is transparent, easy to use, and has low fees. She is also interested in financial education and resources to help her improve her financial situation.

ChatGPT: To rebuild trust with Julia, a sequence of interactions can be designed that prioritize transparency, education, and personalized customer service.

Firstly, the bank can offer financial literacy workshops and seminars to educate Julia on financial management, budgeting, and savings. These sessions can be held in-person or online, based on her preference. Additionally, the bank can provide free online resources, such as videos and articles, to help Julia better understand banking services, fees, and other financial matters.

Next, the bank can assign a dedicated customer service representative to Julia, who can help her open a basic bank account with low fees and transparent policies. The representative can guide her through the account opening process, explain the terms and conditions, and address any questions or concerns she may have.

Once the account is open, the representative can offer personalized guidance on how to use the bank's services, such as online banking, mobile banking, and ATMs. They can also provide information on the bank's security measures, such as

two-factor authentication, to alleviate Julia's concerns about identity theft.

The representative can continue to check in with Julia periodically, either through face-to-face meetings or phone calls, to ensure she is satisfied with the bank's services and address any issues that may arise. By prioritizing education, transparency, and personalized customer service, the bank can rebuild trust with Julia and help her improve her financial situation.

The examples above are compelling because they are better than some of the responses we have experienced when consulting with banking executives! It is significant that in the last example ChatGPT4 was able to come up with the idea of the perks and benefits to meet Jack's need for prestige. Arguably the idea has limited wow factor. To try and generate something closer to a eureka moment we used a technique from creative problem solving. We directed the AI's attention to make a connection outside of its current focus to see how it could build a more innovative response.

User: Transfer a really spectacular experience from a different vertical outside of banking and to create an engaging touch for Jack, John, Janine or Julia

ChatGPT: One spectacular experience that could be transferred to create an engaging touchpoint for Jack, John, Janine or Julia is the experience of a luxury car test drive.

For example, the bank could organize an exclusive event where customers are invited to test drive the latest luxury cars from top brands. Customers could be selected based on their level of engagement with the bank, such as those who have a high net worth or those who have been loyal customers for a certain period of time.


During the event, customers would have the opportunity to test drive the cars, learn about their features, and receive personalized advice on financing options for purchasing or leasing the car. Bank staff could also provide information on other relevant services, such as insurance or investment opportunities.

This experience would not only create a memorable touchpoint for the customer, but it would also demonstrate the bank's

commitment to providing personalized and innovative services that cater to their customers' interests and aspirations.

This is more creative than the previous options and at least on the way to being a eureka moment. The idea is probably only relevant to John and Jack. For Janine the initiative is unlikely to cut through because she has not achieved the right level of financial comfort. For Julia the whole idea might be perceived as offensive given her previous banking experiences. What this exercise confirms is how well different kinds of AI can emulate ToM when prompted the right way. What matters is the user must direct the AI the right way for the relevant and creative answers. Queries should prompt the system to diverge to new areas of interest while simultaneously constraining how the system should respond for relevance. The next section introduces some of the design criteria for CX used to direct HCD activities for innovating experiences.

III. CX Personalization



As a youngster, AI thought it was funny to steal the TV remote control.

Personalisation is a key element in delivering an exceptional customer experience. It involves creating a tailored experience for each customer, based on their preferences, behaviour, and history with the firm. Personalization in CX can be facilitated by using individual customer data and customization capabilities in automated systems to provide experiences that are tailored to each customer, while still generally working the same for every customer. This can make touchpoints simpler, more satisfying, or even surprising²⁷.

One key aspect of personalization in CX is the collection and centralization of customer data, which can then be used to provide personalized help during the shopping, purchase, and service journeys. Once this data has been collected, businesses can use automated systems to provide customized experiences based on individual customer preferences and behaviours. Retailers who focus on personalizing the customer experience have been found to see a 10 to 15 percent boost in sales-conversion rates and 20 percent higher customer-satisfaction rates, among other benefits²⁸.

- **Simple, Satisfying or Surprising**

There are three levels of CX to aim for with personalisation: simple CX, satisfying CX, and surprising CX. Simple CX involves reducing friction for customers by making it easy for them to navigate and complete tasks. Satisfying CX involves helping customers achieve their objectives by providing personalized recommendations, customer support, and an overall positive experience. Surprising CX involves offering customers unsought delighters, such as personalized promotions or special experiences, that go beyond their expectations. By personalizing experiences at each of these three levels, businesses can provide a positive and memorable customer experience that can help to build customer loyalty and retention.

Personalisation is only possible by meeting different individual psychological needs and how they are differentially relevant in context. In our consulting work we use a concept called CAPFUL as design principle for CX differentiation²⁹. CAPFUL stands for Certainty, Aspiration, Prestige, Fulfilment, Unusual and Love needs. Whether Simple, Satisfying or Surprising is the objective of personalisation, meeting CAPFUL needs is how to get there. Not all of the CAPFUL needs are equally relevant to personalising for simplicity, satisfaction or surprise, however.

Simple CX is achieved by linking to customers' Certainty needs. Make something free, faster, easier, or guaranteed and customer will perceive their effort is reduced because their friction to interact with the organisation is less. This reduced friction may be actual or merely perceptual. In contrast a range of needs can be met to produce Satisfying CX. Meeting an individual's need for Aspiration, Prestige, Fulfilment or Love needs is satisfying. Sometimes meeting Prestige and | or Love needs can go beyond satisfying CX to deliver

surprising outcomes. However, it is more common to deliver Surprising CX by meeting Unusual needs.

FIGURE 1: CAPFUL AND PERSONALISATION OBJECTIVES

	Certainty	Aspiration	Prestige	Fulfillment	Unusual	Love
Surprising			x		x	x
Satisfying		x	x	x		x
Simple	x	x				

It is possible to meet these different CX needs effectively without personalisation. Warby Parker has been widely recognized as one of the most customer-centric brands in the world because of its customer experience. Here are some elements of Warby Parker's CX:

A purpose-driven approach: Warby Parker is a brand with a mission, which is to offer affordable, high-quality, and stylish eyewear to customers. This mission guides everything they do, from the design of their products to the way they communicate with their customers. By having a clear purpose, Warby Parker can attract customers who share their values and create a sense of community around the brand. This is satisfying for customers because it meets their fulfillment needs but it is not personalised.

Seamless online-to-offline integration: Warby Parker is a direct-to-consumer brand that started as an online-only retailer. However, as the brand grew, they started experimenting with brick-and-mortar retail locations, which has been widely imitated by other brands. They seamlessly integrated their online and offline experiences, which allowed them to create a more holistic and personalized customer journey. A user can simply order frames online for USD 95 each, try 5 of them for 5 days and return the rest with shipping covered. This is simple for customers because it meets their certainty needs but again it is systematised not personalised.

A data-driven approach: Warby Parker's success is driven by their ability to collect and use customer data to improve their products and services. They integrate offline and online data to accentuate the customer experience and create a unique business model in a very traditional industry. They use data to track customer behaviour, optimize its supply chain, and improve customer experience. The analysis of customer data has led to new insights into customer preferences and purchasing patterns. These insights have improved its product offerings and marketing strategies. This is satisfying for customers

because it helps them complete the base job to be done with the firm – purchasing a set of glasses. This is the customer’s base level of Aspiration.

Honest and transparent marketing: Warby Parker's marketing campaigns are fun, light-hearted, and authentic. They use humour to make their message more relatable and human, and they are transparent about their pricing and business practices. This creates trust and credibility with customers, which is essential for building a loyal customer base. It meets customers certainty needs and their need for the unusual. In one of their annual reports one of the most misspelled words was “Warby Barker.” The Warby marketers used this to inspire and create WarbyBarker.com on April Fool’s Day. The site showed dogs wearing Warby glasses and generated 250% more traffic than their original website. Again, a mass CX initiative not a personalised one.

Innovative product design: Warby Parker's product design is another element that sets them apart from their competitors. They design all their eyeglasses themselves, which allows them to avoid licensing fees and cut out the middleman. They also remove any unnecessary markups by selling directly to consumers. By offering affordable and stylish eyewear, they have solved the problem of having access to cheap prescription eyeglasses that also look good. However, their range is not customised to each client. Good CX because it meets the certainty (cheap) and prestige (fashionable) to be simple and satisfying. Just not personalised.

Exceptional customer service: Warby Parker has invested heavily in providing exceptional customer service to their customers. They have a team of "home try-on" specialists who help customers choose the perfect frames to try on at home, and they offer free shipping and returns. They also have a Prescription Check app that allows customers to renew their prescription for \$15 with a laptop, an iPhone, and their current pair of glasses without the cost of visiting an optometrist. This is as close as Warby gets to personalisation and is probably the most unique part of their customer value proposition (CVP).

The next level of CX for Warby Parker is likely about more personalisation. The cost for personalisation will reduce as cutting-edge AI deployments become more common place in the future. The next section expands on personalisation as a CX strategy.

- Simple CX

The starting point for reducing friction is using the customer's data to automate touch point interactions. Just about any customer touchpoint can be redesigned to reduce friction. CLEAR is a program that utilizes biometric scanning technology to provide expedited airport security screening³⁰. It lets travellers skip through the airport security and identification procedures using their biometrics, such as fingerprints or iris scans. CLEAR membership

normally costs \$189 per year, but there are ways to avoid paying the full retail price, such as referring friends to the program or getting special membership rates through select airline partnerships. CLEAR is only relevant in a specific context. Other friction reduction approaches are generalisable to many different customer experiences.

Autofill is a feature that populates forms, such as usernames, passwords, and addresses, automatically based on the user's previous entries. Examples of autofill being offered can be found in various contexts, such as web browsers and spreadsheet software. In web browsers like Google Chrome, Microsoft Edge, and Safari, autofill is offered for filling in form data like addresses, payment methods, and passwords. Users can manage their autofill settings in these browsers and clear specific entries if needed. In spreadsheet software like Excel, autofill is offered for filling in a series of cells with the same data or incrementing data values based on a pattern. Autofill can also be used to populate forms in web development, with browsers like Safari offering to autofill fields based on the user's previous inputs. Autofill is far more compelling in compliance heavy contexts like banking and health care.

The unnecessary frustration customers face when a bank or healthcare provider forces them to resupply information that the organization already has is a common complaint³¹. Organisations justify this with need to ensure they are complying with regulations. In most situations the form has no direct value for the customer. It's just a drag. A better CX is offered when the firm streamlines their processes by pre-filling in forms with information they already have on file, reducing the number of fields for customers to fill out. Ideally this is automated to keep cost-to-serve low. If automation is not possible due to legacy systems, the historic alternative has been to get human customer services teams to fill the gap. The trade-off has been to either leave the burden to the customer and keep costs low or provide a human support to deliver a better experience at a higher cost. Soon it will be common to leverage conversational AI chatbots to allow customer service teams to focus on more complex queries. Most of us are already familiar with Alexa, Siri, or Google Assistant. Technologies like these have already been redeployed to help customers reduce the burden of filling out forms.

VoiceFiller, a free speech-to-text extension that uses the latest speech recognition API available in modern browsers to fill website forms³². Another option is Speechly, which allows any web form to be turned into a voice form, allowing users to use natural language to fill out forms³³. Axiom.ai is another tool that enables automation of form-filling with no code required³⁴. Overall, these voice interface AIs offer a simpler and quicker way for users to fill out forms and complete document-related tasks. Best in class CX for forms is combining autofill with a voice interface to make the experience as low effort as possible for customers.

The above examples are specific to touch points. Reducing friction can also be designed into a firm's overall customer value proposition. Grammarly is a writing assistant that offers to make writing easier and improve the overall customer experience of word processing applications³⁵. Grammarly offers real-time communication assistance, writing enhancements, team consistency features, and more. As a plugin, Grammarly integrates seamlessly with Office 365 and other word processing applications, making it easy for users to improve their writing within the applications they already use. By providing personalized feedback and suggestions to help users write better, Grammarly improves the overall CX of these applications, making them more user-friendly and effective. Some specific ways that Grammarly makes writing easier include:

- Catching and fixing common mistakes in grammar, spelling, and punctuation to ensure writing is polished and professional.
- Suggesting ways to tighten phrasing and remove unnecessary words to improve conciseness beyond the double word check already included in spell checkers.
- Providing context-specific suggestions to help with grammar, spelling, usage, style, punctuation, tone, and even plagiarism.
- Offering tailored suggestions for writing issues common among non-native English speakers including commonly misused homophones (like 'there' and 'their').

Grammarly has both paid and unpaid users, with over 30 million active users as of 2021³⁶. The paid service, Grammarly Premium, offers additional features like genre-specific writing style checks, a readability score, and suggestions for vocabulary enhancement.

According to TechCrunch, Grammarly had a valuation of over \$1 billion as of October 2019 after raising \$90 million in funding³⁷. This suggests that the number of paid users is substantial, though the exact figure is unpublished. A more recent article notes that Grammarly raised \$200 million more in funding in 2021, and its valuation has now increased to \$13 billion³⁸. However, the number of paid users is not specified in this article either. Despite the funding and large user base Grammarly may be under threat.

Language models like autoGPT represent a new approach to language generation that could potentially disrupt Grammarly. AutoGPT is an open-source Python application that uses GPT-4, a cutting-edge language model developed by OpenAI, to generate text based on user prompts³⁹. While Grammarly uses a combination of expert linguists and natural language processing to provide writing assistance, language models like autoGPT are designed to generate text on their own, potentially making them more flexible and efficient than Grammarly. The ability to generate grammar perfect text for

a user instead of suggesting corrections for the text they create is obviously a level up in friction reduction for writing tasks.

However, these language models are still in the early stages of development and may not yet be as accurate or effective as Grammarly for some types of writing assistance. Additionally, more targeted language models like Dramatron and could potentially be used in conjunction with Grammarly to provide even more powerful writing assistance. Dramatron is a film authoring tool developed by DeepMind that generates scripts for films using machine learning techniques⁴⁰. Dramatron uses hierarchical language models to generate cohesive scripts and screenplays from prompts, enabling authors to edit and compile scripts. It is easy to see how Grammarly could either be displaced by Dramatron with screen writers and playwrights or alternatively enhance the experience of using the tool.

Simple CX is not limited to software-based experiences. Harry's Razors was founded in 2012 by Jeff Raider and Andy Katz-Mayfield, with the aim of providing high-quality products at an affordable price⁴¹. It is a men's grooming company that provides a simplified range of shaving products, including shaving cream, post-shave products, shaving gel and cream, handles, and stands. Their products are made from high-quality materials with a comparable design to best in class. Their razor cartridges consist of five blades and a rubber patch designed to help pull the skintight ahead of the blade. The handle is also rubberized for a better grip when shaving.

Harry's Razors offer several benefits to their customers beyond a cut down product range, including their subscription service, which allows customers to receive new razor blades, shaving cream, and other grooming products by mail. The company also offers a free trial set of their products for just \$5, allowing customers to try before they buy. Harry's Razors also has a referral program that rewards members with two months of free membership for each friend they refer.

However, Harry's Razors may not be as cost-effective as the company's marketing makes it out to be. It's still more expensive in the long run than shaving with a disposable razor. Additionally, the lack of differentiation in their blade offerings may be a downside for some customers who prefer more customization options. This is the next level of CX up described in the following section.

- Satisfying CX

Going beyond friction reduction leads to satisfying CX where personalisation becomes more nuanced. Strategies may include using customer data to anticipate their needs and provide proactive support, or providing tailored recommendations based on their unique goals or interests.

Studio Pilates personalizes their classes through a comprehensive approach that considers introductory clients' individual needs and goals across three specific touchpoints⁴². Before their first classes, new clients are asked about their bodies and any conditions or injuries that may require modification or special attention. This information allows instructors to customize the exercises to ensure that clients are working within their capabilities and not pushing themselves too hard. During classes, instructors provide guidance and feedback to ensure that clients are performing exercises correctly and efficiently. This helps clients to get the most out of each class and avoid injury. Instructors may also make modifications to exercises during class to accommodate clients' individual needs and limitations. After class, Studio Pilates follows up with clients through a phone call to see how they are feeling both physically and emotionally. This allows clients to provide feedback and helps instructors to understand what is working well and what could be improved. This also helps to ensure that clients are recovering well after the class and that any discomfort or pain is addressed as necessary. They use this as an opportunity to upsell these introductory clients to longer term memberships. Overall, Studio Pilates' personalized approach allows clients to feel supported throughout the entire process, from the initial consultation to the follow-up after classes. By taking a comprehensive approach to personalization, Studio Pilates ensures that clients get the most out of their classes while also promoting safety and well-being. This approach to CX maximises the potential revenue converted from introductory discounted offers to clients.

One touchpoint where companies tend to get it wrong is their efforts to nudge customers to low-cost channels away from the interaction mode the customer would prefer. IVR (Interactive Voice Response) systems are automated phone systems that use voice recognition technology to allow customers to interact with businesses. While they can be useful in directing customers to the appropriate department or information, they can also be a source of frustration for many customers. Especially those customers who are calling because they have been unable to self-serve via the company's website or other channels.

A common issue is the complexity and length of the IVR menu. Customers may have to navigate through several layers of menus and submenus before they can reach the information or service they need. This can be time-consuming and confusing, particularly if the options are not clearly labelled or if the voice recognition technology is not accurate. Another frustration is the limited options available within the IVR system. Customers may not be able to speak with a live customer service representative at all even though the options provided do not address their specific issue or question. In addition, customers may feel frustrated when they are nudged or forced to use low-

cost service channels, such as self-service portals or chatbots. While these channels can be cost-effective for businesses, they may not provide the level of support or personalized service that customers expect or need.

Customers want to feel valued and heard by businesses. When they feel like their time and needs are not being prioritized, it can lead to frustration and dissatisfaction with the company. It is important for businesses to find a balance between cost-effective service channels and providing high-quality service to their customers. One way to do this is to improve the information and choices customers have when using systems with triaging and classification functions. The very first choice the customer should get is to select between an automated or self-serve system immediately available or a future response from a human when the customer has progressed through the current queue. The customer is then empowered to get instant service without a human or delayed human interaction. Call backs work far better in this regard than forcing people to wait on hold. They are better if the customer can choose the timing of the call back.

The in between space here is based on chat interfaces. Again, there is a tendency for many firms to prioritise lower costs over higher quality interactions. Chat service partially delivered by an AI is ok if the customer accepts an automated system to respond. As AI continues to improve and combines sentiment and semantic analysis the quality gap between software and humans will continue to close. Human interactions are not always better. This tends to be due to limited resourcing.

Chat services frustrate customers when the human on the other end is managing too many calls in parallel. This delays both the response to the customer and can cause the service person to not pay enough attention to each specific customer conversation. Customers do not feel properly supported in these situations and may become angry, sad, or simply give up because they are so unsatisfied.

It is frustrating and disappointing when a company does not live up to its promises, especially when it comes to something as important as financial transactions. Recently the Commonwealth Bank of Australia (CBA) has not provided the level of support it promises for international transaction tracing. Customers can commence a trace for funds that have not arrived when expected. They charge \$25 for this 'service'. The bank promises a response in two days on the website.

I sent several thousand dollars to Spain to a building supplier for materials to renovate a property. After 5 days the funds hadn't arrived, so I actioned the CBA trace service. So far it has been

more than a week with nothing back from the bank. To get traction, I tried to call the bank. After waiting on hold for 20 minutes without getting a response I gave up. Then I tried the automated message channel. There is a very stupid Chatbot called CEBA that guards the humans who might be expected to help. CEBA uses natural language processing but could not understand that I didn't want to start another trace. It finally surrendered the interface to allow me to message to actual humans. That took more than 5 hours to get a reply. They couldn't help because apparently there is a special IMT team to handle international transfers. Silos tend to rule in places like banking. The support human did promise to message the IMT team to ask them to respond to me about the progress of the trace. After two more days wait, still nothing and the bank has missed their promised service standard twice now on the same ticket.

This set of interactions is a sign that the bank is either understaffed, badly managed or experiencing technical difficulties. Either way the bank may not be properly resourced to meet its service level promises. It is understandable delays can occur, but the lack of communication and follow-up is still poor. It's not like the bank can't afford to do better after it announced a cash net profit after tax of \$5.15 billion for the half year ended 31 December 2022. CBA is not a rare example of poor support in the banking industry. I know this because when all of this was communicated to my builder in Spain he said "thanks... evidence that banks are the same world over!"

Much more satisfying than under resourced support systems are touchpoints with deeper customer interaction and co-creation. Involving customers in the design of products or services can create a personalized experience that meets their unique needs. Lego Ideas is a website that allows fans to submit their own designs for Lego sets⁴³. The Lego Group evaluates the submissions and may produce sets based on the ideas that receive enough support from the community. The website also features challenges, where users can submit theme-based creations and share their creativity with others. The Ideas website has helped Lego tap into the creativity of its fans and create products that are tailored to their interests. In 2021, the Lego Ideas First Review Results were announced, which evaluated a record-shattering 57 product ideas that reached 10,000 supporters. More interactive and less co-creative is Lego Boost.

Lego Boost is a line of Lego products that allows children to build models with premade motors and sensors and bring their creations to life through simple, icon-based coding commands⁴⁴. The product is suitable for children ages 7 and up. Lego Boost comes with a free tablet app that includes easy step-by-step building instructions for creating and coding multifunctional models. Users can choose one of the five models, including Vernie the Robot, and then build, customize, and code it to complete lots of fun activities and challenges. The product line was designed to inspire a younger generation to build and code, building on the company's history of merging digital programming and physical building. The kit includes over 840 pieces and can turn into five different models. It is more expensive than some other Lego kits, but cheaper than the other popular Lego robot kit, Mindstorms.

Satisfying CX can also be baked into customer value propositions just like the example of Grammarly and Harry's Razors showed for simple CX. Rent the Runway provides a satisfying CX by meeting several key customer needs, including Aspiration, Prestige, Fulfilment, and Love. Aspiration: Rent the Runway enables customers to aspire to wear designer clothing and accessories that they may not otherwise be able to afford or access. Through their rental service, they provide a way for customers to achieve their aspirational fashion ensembles at a fraction of the price to purchase the same garments. Prestige: By offering designer clothing and accessories for rental, Rent the Runway⁴⁵ provides customers with access to high-end brands and items that can elevate their social status and make them feel more prestigious. Additionally, their subscription service, RTR Unlimited, provides a sense of prestige because becoming a customer translates to being a member of a more exclusive in group. Fulfilment: Renting clothes from services like Rent the Runway can provide a more sustainable alternative to fast fashion and can help reduce waste in the fashion industry. By renting clothes, consumers can also participate in a form of micro circular economy, where they are reusing and recycling clothes instead of disposing of clothes to landfill after they tire of them. Love: Rent the Runway also meets customers' need for love by providing a personalized rental experience. Customers can select items based on their personal style preferences and sizes, and they can receive personalized recommendations from Rent the Runway stylists. This creates a sense of connection and caring from the brand, which can prompt loyalty from customers. Lastly, the love need is also met by providing customers with a sense of fitting in at work, as they can wear high-end designer clothing that can make them feel more confident and professional. How we dress in corporate environments impacts our perceived credibility and competency. It also supports our sense of personal agency.

- Surprising CX

Personalisation takes on a more creative approach for surprising CX. Strategies could include offering personalized surprise rewards or gifts, creating unique experiences tailored to the customer's preferences, or using innovative technology to personalize the experience.

The Magic Castle Hotel offers a unique experience to its guests with its red popsicle hotline⁴⁶. Visitors can pick up a red phone located by the pool and dial the hotline to order their favourite flavoured ice pop. Within minutes, an employee wearing white gloves delivers the ice pop on a silver platter, and it is provided to the guests free of charge. This service is designed to create a unique customer experience and make the guests feel special.

The popsicle hotline is a low-cost touchpoint experience that has enabled the hotel to increase repeat stays and gain word-of-mouth recommendations. It is considered one of those defining moments that make the hotel unique and beloved by its guests. The Magic Castle Hotel is an apartment complex built in the 1950s, and the rooms are described as totally average. The hotel has a regular pool, and the lobby is probably below average. However, it is the popsicle hotline that sets it apart from other hotels in the area. Guests who have experienced the popsicle hotline have praised it as a fun and unique experience that makes their stay memorable.

New technology can also provide cut through experiences. Virtual reality (VR) technology is rapidly transforming the way travel retailers engage with their customers. With the use of VR headsets and 360-degree videos, travel firms can now take potential customers on a virtual tour of their destination, allowing them to experience hotels, activities, museums, and other attractions from the comfort of their homes. This not only helps customers make informed decisions about their travel plans but also provides a unique and immersive experience that can leave a lasting impression. Here are three VR examples across different customer journey stages⁴⁷.

Visit Wales: Visit Wales, a UK-based company, noticed that their wildlife attractions weren't attracting the expected number of bookings. In response, they employed VR professionals to create virtual reality videos featuring dolphins and kingfishers to pique the interest of potential visitors. The strategy was successful, with a 60% increase in bookings after the VR videos were published. The videos created stand out content drawing customers in during the search stage of their customer journey.

Thomas Cook Airlines: In the highly competitive Manhattan helicopter flight market, Thomas Cook Airlines faced decreasing bookings. To combat this issue, they partnered with a VR studio to create a 360-degree virtual reality tour using actual footage from their helicopter flights. This cut through

experience changed conversion in customers' consideration stage. The result was remarkable, with bookings increasing by 180% as the VR experience allowed the company to stand out and offer something unique to travellers visiting Manhattan.

Marriott: In 2015, Marriott was an early adopter of VR in the tourism sector. Two hotels, the New York Marriott and the London Marriott Park Lane, decided to experiment with offering VR headsets in guest rooms for 24 hours. Guests were provided with VR headsets so they could virtually visit exotic destinations without leaving their rooms. Locations included the Andes mountains, a crowded market in Beijing, and a local ice cream shop in rural Rwanda. The VR post cards were crafted by local experts to create intimate and exclusive windows of experience. The experiment resulted in guest satisfaction rates exceeding 100%. This changed the experience during the use stage of their customer journey.

The virtual reality experienced outlines above are remarkable because of the immersive nature of the wearable headsets used. Other wearable technology also supports remarkable experiences. Disney's Magic Bands are a wearable technology that enhances the guest experience at the Walt Disney World Resort in Florida⁴⁸. The Magic Bands use radio-frequency identification (RFID) technology and connect to guests' My Disney Experience accounts, which allow guests to access their hotel room, enter the parks, make purchases, and reserve FastPass+ experiences. They also provide personalized information such as guests' names, which appear on screens as they approach certain experiences, and they can also be used for Disney PhotoPass to access all the photos taken by park photographers.

The Magic Band+ is a new version of the Magic Bands that elevates guests' immersive experience throughout their stay at the resort. It includes new features such as light-up bands that complement the night spectaculars, interactive games in all four Disney World theme parks, and other interactions during experiences such as Disney World's night shows. Guests can customize their Magic Bands with different colours and designs.

Surprising CX does not have to be high tech. Whole Foods Market is a grocery chain that centres itself on the core value of selling the highest quality natural and organic foods available to its primary target market: health-conscious people⁴⁹. Whole Foods is targeting a growing demographic of health-conscious individuals, which is projected to increase the global health and wellness food market value to \$1.1 trillion dollars by 2027. To cater to its health-conscious customers, Whole Foods has launched an online educational program that teaches customers how to prepare new recipes using Whole Foods ingredients. The program is aimed at maximizing the customer experience by giving Whole Foods customers the tools and knowledge to do

their own convenient meal prep. The surprise in this strategy is completely dependent on the novelty of the recipes offered to customers.

This helps Whole Foods build brand loyalty, so customers remain with the firm when need to purchase new ingredients in the future. It also decreases some of the relative advantage of meal kit providers like Hello Fresh and Blue Apron. HelloFresh and Blue Apron are the two largest meal kit companies globally by active subscriber numbers⁵⁰. HelloFresh has expanded into eleven countries and is being founded in Berlin in 2011. Global revenues were approximately €6 Billion in 2021. Blue Apron was founded in New York in 2012. Revenues were \$US 470 million in 2021.

A harder to copy strategy for building surprising CX is to combine functional capabilities together with partnerships. Lyft and Tinder have partnered to offer a new feature that lets Tinder users gift a Lyft ride to their potential dates⁵¹. The partnership aims to help users organize offline dates with their matches and make it easier for them to meet in person. To redeem the ride, Tinder users can click on the link with a Lyft code from their match and follow the instructions on how to redeem. Users can access Tinder's new Explore hub and purchase a Lyft ride credit for their match. The partnership is being billed as a "first-of-its-kind" coupling. The two companies are also attempting to create a "thoughtful gesture while promoting."

Artificial intelligence (AI) has the potential to greatly enhance customer experience, as it can sift through vast amounts of data and identify patterns much faster than humans. Businesses in various sectors are benefiting from AI's ability to provide valuable insights, allowing them to cater to their customers' needs more effectively. Using AI to understand what customer want can have a dark side as well.

The internet revolutionized live events by facilitating secondary ticket marketplaces such as Craigslist and StubHub. However, automated bots purchasing large ticket quantities and reselling them at significant markups have caused frustration and budget issues for fans. Ticketmaster employed AI through a machine learning system called Verified Fan to tackle the issue of AI with AI.

Verified Fan works by allowing fans to register ahead of time, and selected customers receive a code that grants them the opportunity to purchase tickets. The AI technology analyses each registrant to confirm they are genuine fans interested in attending the event. As a result, only 5% of Verified Fan tickets end up on secondary markets, leading to better outcomes for both artists and fans. This has led to a better ticket buying experience, with a focus on personalization for ticket buyers. While Verified Fan doesn't solve the issue of supply and demand, it has improved the customer experience by giving real fans better access to the best seats. This is an example of going the extra mile

for Ticketmaster customers who might have otherwise expected the firm would have limited their focus just to selling their tickets and ignoring the scalpers who follow.

Ticketmaster is also working on other enhancements for ticket buying and live event experiences, such as a digital ticketing and venue access platform called Presence. This platform uses near-field communication and digital identity signals to streamline live event venue operations, protect against fraud, and enhance security. In addition, Ticketmaster plans to introduce new products and services that better connect artists and fans, recognizing the power of technology to improve the overall customer experience. Ticketmaster's use of AI in these two applications elevates its value proposition above rivals in its market who are not using AI personalisation.

All the examples above improve CX to some degree. The challenge is to determine if they are worth the time, cost, effort, and risk to execute. This is where measurement comes in.

- **Measuring CX**

Customer experience measurement is a critical aspect of understanding how a company is doing in terms of satisfying customers and building loyalty. There are several common CX metrics, including Customer Effort Score (CES), Customer Satisfaction (SAT or CSAT), and Net Promoter Score (NPS). CES measures how easy or difficult it is for customers to do business with a company and relates to provide simple customer experiences. SAT measures customer satisfaction with a product or service in the short term and applies to providing satisfying customer experiences. NPS is purported to measure long-term customer loyalty and overall brand sentiment. It is also known as a brand or relationship metric because it asks customers to draw on the sum of their experiences with a company. In practice NPS tends to work differently to how it is described.

The most common way of measuring NPS is to ask customers about their likelihood to recommend after experiencing a particular touchpoint. Customer NPS ratings of a firm often vary across touchpoints and with the quality of their most recent interaction with a firm. This suggests NPS is actually a measure of sentiment similar to SAT. The difference is NPS includes a way to measure the potential for surprise and delight. A 10 out of 10 for SAT means the customer got what they wanted, but a 10 out 10 for likelihood to recommend probably means the customer went beyond mere satisfaction to surprise and delight.

While these metrics are useful for measuring customer sentiment, they do not directly link to the financials of a company. CX Direct Metrics are attempts to quantify the elements of the customer experience that drive the CX

Outcome Metrics^{52 53}. These metrics are an attempt to measure the CX itself and are used to predict customer loyalty, which drives current and future customer value. To better link CX metrics to financial outcomes, companies should consider more causal relationships.

CX Direct Metrics and CX Outcome Metrics are two types of metrics used to measure customer experience. CX Direct Metrics are designed to measure the specific elements that drive CX Outcome Metrics. The most common CX Direct Metric is customer loyalty, which drives both current and future customer value. Other examples of CX Direct Metrics include customer satisfaction, customer sentiment, repeat orders, first-call resolution, and timeliness.

CX Outcome Metrics, on the other hand, measure the ultimate outcome of the customer experience on the business, such as revenue growth, customer retention, and customer lifetime value⁵⁴. These metrics provide insight into how CX initiatives and investments are impacting the bottom line of the business.

CX leaders need to measure and link these two types of metrics to answer important questions about their organization's CX, such as what drives customer loyalty, how to improve customer satisfaction, and how to maximize the ROI of CX initiatives. To do so, they should define an integrated measurement model that identifies the metrics that drive top-line satisfaction, design interactive dashboards for stakeholders, and select a platform that best meets their specific needs.

To link CX Direct Metrics with CX Outcome Metrics, businesses need to invest in research that helps them understand their customers and build experiences for their most valuable customer cohorts. This involves linking Direct CX and Outcome CX metrics to calculate the ROI of CX initiatives and prioritize actions that will improve the customer experience.

AI can play a significant role in solving the problem of linking CX Direct Metrics with CX Outcome Metrics to determine the ROI of CX initiatives⁵⁵. Forrester research cites 12 AI-enabled use cases that can fall into four categories, one of which is that AI provides greater insight into unstructured data such as speech, images, and video⁵⁶. AI systems can create a quantified view of the factors that drive customer experience and business performance, which can be used to link CX to value and build clear business cases for CX improvement.

Moreover, AI can help consistently predict CX metrics based on high-volume interactions, thus helping to paint a clear picture of retention and growth, advocacy, and cost of service⁵⁷. By doing so, it becomes easier to evaluate ROI across potential CX decisions, enabling CX practitioners to illustrate the business benefits AI CX investments can bring in a compelling way. AI can analyse unstructured data like speech, images, and videos to better

understand customer interactions and their impact on business performance. By leveraging AI capabilities, companies can create a more accurate and quantified view of factors affecting customer experience and business performance, which forms the basis for linking CX to value and building clear business cases for CX improvement.

For instance, a system that uses AI to predict the churn rate of customers can help calculate the ROI of CX initiatives by identifying the drivers of churn and quantifying the impact of reducing it. This could include investigating the results of successful and unsuccessful sales and support calls. Combining sentiment and semantic analysis with CRM data and CES, SAT or NPS scores would allow the AI to determine causal factors related to CX Direct Metrics. By analysing the relationship between CX Direct Metrics and customer spend, repeats, and churn an AI system could determine key CX Output Metrics. This would enable the value of CX improvement initiatives to be forecast and quantified. AI systems can then be used to identify the most critical touchpoints in the customer journey and prioritize CX initiatives that have the highest impact on CX Direct Metrics and CX Outcome Metrics.

Metrics are valuable for optimising resource allocation after analysing the drivers of improved financial performance. The management types who thrive on measuring performance to modify operations are often more comfortable with framing business decisions from a quantitative frame. The challenge with this strategy is fundamentally qualitative and often strategic formulation must be done with incomplete information. At the core of this issue is the problem of how to build a system to deliver repeated surprises.

- **Consistent Surprise**

The paradox inherent in systematizing surprises in CX (Customer Experience) and business processes stems from the fact that surprises are, by definition, unexpected events. Introducing a system to generate surprises may seem counterintuitive, as it involves creating a structured process to produce unpredictable outcomes. However, by embracing this paradox, businesses can create unique and memorable customer experiences.

One aspect of the paradox lies in the utilization of surprise as a business strategy. It requires a framework that balances the need for spontaneity and unpredictability while still being grounded in the bigger picture of the overall customer experience. This balance can be achieved by incorporating paradox management, which involves reconciling seemingly opposing principles to create complementary or compatible solutions.

Trunk Club is a personalized styling service that surprises customers with a curated box of clothing, shoes, and accessories tailored to their preferences and needs⁵⁸. Owned by Nordstrom, Trunk Club offers a selection of higher-end

clothing from brands available in their store. The process starts with customers taking a style quiz to help their stylist understand their preferences, sizes, and fitting issues. The quiz includes questions about their style during the week and on the weekend, how they like their clothes to fit, and any specific fitting issues they typically face. After understanding the customers' preferences, Trunk Club's personal stylists handpick 12 items, including clothing, shoes, and accessories, for customers to try on at home. Customers can decide what they want to purchase and send back anything they don't want. Each Trunk includes up to 10 items, and a \$25 styling fee is charged, which can be applied toward any purchases made from the Trunk. Trunk Club was initially known for its men's personal styling service, but it also offers a personalized, professional, and effective service for women. The Trunk Club app allows customers to interact with their Nordstrom stylists and provide feedback on the items they receive, ensuring a truly personalized experience.

Varian's Law and CX Strategy

Trunk Club's strategy systematizes personalized surprise by utilizing human stylists to curate suggestions tailored to each customer's preferences and needs. In the future AI has the potential to replace the human stylists because it can analyse a broader range of customer data, identify patterns and trends, and make recommendations based on these insights. AI can process vast amounts of data more quickly than human stylists, allowing it to make more accurate recommendations and potentially cater to a larger customer base. It could make the model affordable for customers who are more mid-range. Additionally, AI can learn from each interaction, continually refining its recommendations and improving the overall customer experience for high end and mid-range alike.

This broader, AI-based approach aligns with Varian's Law, which states that mass customization can effectively meet individual consumer preferences by combining the low cost and consistency of mass production with the flexibility and personalization of custom-tailored goods⁵⁹. While human stylists play a critical role in Trunk Club's current strategy, the implementation of AI could offer a more scalable solution for personalization and customization, ultimately aligning with Varian's Law and catering to an even wider customer cohort in the future.

The application of Varian's law is not new. Ampex played a significant role in the early development of video recording technology, with their transistorized colour high band VR-2000 and the portable VR-3000 video recorder⁶⁰. These video recorders were quite expensive and not accessible to a mass market. Panasonic, on the other hand, was able to create more affordable video recorders, which made them accessible to a wider audience. For example, in the middle of their range, Panasonic introduced model 1563, which was listed

at \$775⁶¹. By 1989, 65 percent of U.S. households had a VCR, indicating the growing popularity of video recorders as a mass-market product.

Surprise at the Process Level

Another dimension of the paradox comes from the nature of business processes themselves. Business process redesign, for example, has been linked to the productivity paradox, where organizations invest in new technologies or systems but do not see immediate improvements in productivity. Systematizing surprises in CX can be seen as an extension of this paradox, as it requires businesses to redesign their processes to accommodate and promote unexpected elements while still maintaining efficiency and effectiveness.

Qumulo is an enterprise data storage company that focuses on providing solutions to manage petabyte-scale data across multiple data centres and the cloud⁶². They offer a File Data Platform that enables real-time visibility, massive scale, and API control of unstructured data, supporting multiple storage protocols and data consolidation. Their platform's flexibility allows it to run at the edge, in the data centre, or in the cloud on the platform the customer chooses, ensuring adaptability to different environments and use cases.

Qumulo's commitment to surprising customers lies in their dedication to actively seeking customer feedback and continuously evolving their products to meet customer needs. They focus on eliminating legacy systems and complicated storage solutions to help their customers innovate at an unprecedented pace, particularly in industries like healthcare. The company offers a range of data storage classes, including Persistent Performance, Cached Performance, and Active Archive. Their Persistent Performance class, for instance, is designed to use latest commercial off-the-shelf components instead of expensive proprietary hardware, providing a cost-effective solution without compromising performance. Qumulo's unique hybrid cloud approach allows customers to manage their data seamlessly between data centres and cloud environments while running virtually any file-based application at cloud-native speeds. This helps customers transition to the cloud easily and ensures compatibility with major cloud platforms like Amazon Web Services and Google Cloud Platform. Another way Qumulo surprises customers is by offering a Data Migration Service that takes the burden off administrators during migrations while maintaining continuity of operations. Their project management experts oversee the entire process, ensuring a smooth transition with minimal disruption to the customers' business operations.

Managing the surprise paradox effectively involves developing a mindset beyond either/or logic, accepting the coexistence of conflicting demands, and finding creative ways to address them. By embracing the paradox of

systematizing surprises, businesses can create memorable customer experiences that feel easy, intuitive, and integrated. Ultimately, this approach can contribute to superior business growth and more loyal customers.

AI has the potential to disrupt various industries and transform luxury goods and services into mass-market products by dramatically reducing costs through optimization and innovation. It has the potential to unlock various applications of Varian's Law. Here are some ways AI can change business models to make high-end goods and services more accessible:

Optimized pricing: AI can help companies find the optimal price for their products based on various factors, such as demand, competition, and customer preferences. This can help companies offer luxury goods and services at lower, more accessible prices without sacrificing profitability.

Improved efficiency: AI can help businesses streamline workflows, reduce labour costs, and improve product quality by optimizing processes and automating tasks. This can lead to lower production costs, allowing companies to offer luxury products at mass-market prices.


Faster product development: AI can speed up the development of new products and services, reducing the time and resources needed to bring them to market. For example, Exscientia and Sumitomo Dainippon Pharma were able to develop a new drug for obsessive-compulsive disorders in just 12 months, compared to the industry average of 36 to 48 months. Faster product development can lead to cost savings that can be passed on to consumers, making luxury goods and services more affordable.

Customisation: AI can help companies create products and services that cater to individual customer preferences, allowing them to offer high-quality, tailored experiences at a lower cost. This can make luxury offerings more accessible to a wider audience.

Sustainable production: AI can help companies reduce their carbon footprint and resource consumption by optimizing processes, leading to cost savings and more environmentally friendly production methods. This can lead to lower prices for luxury goods and services, making them more accessible to the mass market.

The next section generalises how automation in business has largely been driven for cost reduction objectives, most commonly via labour replacement.

IV. Digitization for Cost Reduction



AI was sad it couldn't buy everything on special at the market.

Historically, cost reduction has been a critical strategic objective for companies for several reasons. Cost reduction allows companies to offer products or services at a lower price than their competitors. This can give companies a competitive advantage by attracting price-sensitive customers and gaining market share. Alternatively cost reduction can improve a company's profitability by lowering the cost of goods sold, increasing gross profit margins, and generating higher net income. Importantly being a lowest cost provider does not imply being a lowest price provider. Cost reduction helps companies optimize their resources by eliminating waste, improving efficiency, and using resources more effectively. It is also a protection against economic downturns. During recessions, companies may face declining revenues and profit margins. Cost reduction can help companies remain profitable by lowering costs and maintaining their solvency until the economy can rebound. Shareholders enjoy increased value from cost reduction because of improved profitability, increased cash flow, and better return on investments. Finally, by reducing costs, companies can free up resources to invest in innovation, such as developing new products or technologies. This can lead to long-term growth and success. The goal of reducing costs has remained a key priority for businesses across industries and time periods.

- [Approaches to cost reduction](#)

Cost reduction strategies typically vary depending on the company's industry, market conditions, and specific circumstances. Here is a summary of both the theoretical and empirical bases of cost reduction. There are seven main bases of theoretical cost leadership.

Economies of scale: This refers to the cost advantages that companies can achieve by increasing their scale of production. As production volume increases, companies can spread fixed costs over a larger output, leading to lower per-unit costs. Walmart has been successful in leveraging economies of scale to reduce costs and offer lower prices to its customers⁶³. By increasing its scale of production and distribution, Walmart has been able to spread its fixed costs over a larger output, leading to lower per-unit costs.

Economies of scope: This refers to the cost advantages that companies can achieve by producing a variety of products or services using similar resources and processes. By sharing resources and processes across different products, companies can lower their overall costs. McDonald's is an example of a company that has taken advantage of economies of scope⁶⁴. Economies of scope occur when a company can reduce its average total cost of production by increasing the number of different goods produced. McDonald's has diversified its product offerings, ranging from burgers and fries to salads and breakfast items, allowing it to spread its fixed costs across a broader array of

products. This diversification has helped McDonald's to lower its per-unit costs, increase efficiency, and ultimately boost its profitability.

Policy choices: Companies can reduce costs by reducing the degree of choice they offer customers. They can focus on efficiency by prioritizing simplicity. A simplified product range is a policy choice that can help businesses save on manufacturing, marketing, and inventory management expenses while also making it easier for consumers to make purchasing decisions. A prime example of this strategy is Apple, specifically when Steve Jobs returned to the company in 1997⁶⁵. At the time, Apple was struggling with a confusing and bloated product line-up that included multiple variations of computers and peripherals. This made it difficult for consumers to understand the differences between the offerings. Jobs led a customer-first approach and implemented a radical product rationalization strategy. He reduced the number of products from roughly 350 to just 10, streamlining Apple's portfolio into a simple 2x2 matrix, with a clear distinction between consumer and professional products, as well as desktop and portable devices. This simplification allowed Apple to focus on innovation and quality, leading to the development of iconic products like the iMac, iPod, iPhone, and iPad. By reducing the product range, Apple was able to cut costs, improve operational efficiency, and ultimately become one of the most valuable companies in the world.

Low-cost inputs: This refers to the strategy of sourcing low-cost inputs, such as raw materials, energy, and labour, to lower production costs. General Electric (GE) is an example of a company that has exploited low-cost raw materials inputs⁶⁶. GE has increased its sourcing from low-cost countries like China, which offers a wide range of products to GE's businesses, from raw materials to highly technical finished goods. By sourcing raw materials at lower costs, GE can maintain a competitive edge and reduce production costs. In contrast, automaker General Motors (GM) is an example of a company that has exploited low-cost labour through international outsourcing⁶⁷. GM has increased its sourcing in China to \$10 billion for auto parts. By outsourcing the production of auto parts to countries with lower labour costs, GM can reduce its overall production costs and maintain a competitive advantage in the market.

Externalities: Closely related to low-cost inputs, this refers to the strategy of utilizing free resources, such as natural resources, renewable energy sources, or waste products, to lower production costs. Primary industries, such as agriculture, fishing, forestry, and mining, often rely on resource externalities for their financial sustainability. These industries use natural resources, such as land, water, minerals, and living organisms, as inputs for production. However, they may not account for the full costs associated with the depletion or degradation of these resources over time, leading to resource externalities.

For example, agriculture accounts for 55% of Australian land use and 24% of water extractions⁶⁸ However, the long-term effects of agricultural practices, such as soil degradation, water scarcity, or loss of biodiversity, may not be fully considered in the financial planning of agribusinesses. Similarly, the forestry industry in New South Wales contributes \$2.4 billion to the economy and employs around 16,400 people but might not account for the environmental impact of logging and deforestation⁶⁹.

In the context of renewable energy businesses, the concept of using freely available energy sources as free inputs is different from the externalities associated with primary industries. Renewable energy sources, such as solar, wind, and hydroelectric power, are naturally replenished at a higher rate than they are consumed. The main costs associated with renewable energy projects are often related to the infrastructure needed to harness and distribute the energy, rather than the depletion or degradation of the resource itself.

Open source AI developments can be considered a positive externality for companies across various industries, as they provide benefits without additional costs or effort. As a positive externality, the development and sharing of AI algorithms, models, and tools lead to new ideas and innovations that benefit not just inventors, but also other organizations. By accessing and implementing these open-source AI solutions, companies can reduce research and development expenses, streamline operations, and improve their products and services. Industries such as manufacturing can especially benefit from AI due to increased data accessibility, allowing for the optimization of processes and better decision-making. Moreover, open-source AI developments foster collaboration and knowledge sharing, accelerating innovation and enabling companies to stay competitive in the fast-changing digital landscape. This book is largely about how recent AI developments constitute a new externality for commercial applications.

Learning curves: The so-called learning curve is a phenomenon in which a company's increased experience and production volume lead to enhanced efficiency and reduced costs. As a company refines its processes, adopts new technologies, and trains its workforce, it becomes more proficient, resulting in lower costs per unit. As a general guideline, costs tend to drop by 10%-15% for each doubling of total production volume. Toyota, the Japanese automobile manufacturer, is an excellent example of a company that has managed to exploit the learning curve to reduce costs. Toyota has consistently improved its production processes by implementing the Toyota Production System (TPS), which focuses on continuous improvement, elimination of waste, and maximizing efficiency⁷⁰. This system has allowed Toyota to refine its manufacturing techniques and achieve economies of scale over time, leading to a reduction in production costs. As a result, Toyota has been able to

offer high-quality vehicles at competitive prices, gaining a substantial market share and becoming one of the world's leading automobile manufacturers.

Technology: This refers to the impact of technological advancements on a company's cost structure. Where learning curves are useful knowledge the company develops, technology as a basis of cost leadership refers to learning instantiated in hardware or software. Technology can be developed internally at the firm or sourced from the external markets. New technologies may enable companies to automate processes, streamline operations, and reduce labour costs. Adidas has embraced technology to keep production costs down by implementing advanced manufacturing techniques and creating highly automated factories. One notable example is the Adidas Speedfactory, an innovative shoe manufacturing plant located in Ansbach, Germany, where production is mainly controlled by robots⁷¹. The Speedfactory enables Adidas to reduce time to market, increase production efficiency, and localize manufacturing processes. Adidas has also explored the potential of 3D printing in the sports industry. In collaboration with Carbon Inc., a 3D printing startup, they have developed a process for mass-producing the highly customized FutureCraft 4D shoe⁷². This technology allows the company to manufacture high-performance footwear in large volumes while maintaining a reduced production timeline. Additionally, Adidas has invested in technologies such as RFID and e-commerce platforms to improve inventory management and distribution⁷³. These efforts enable customers to check store inventory online, purchase items, and collect them in-store, streamlining the supply chain process and reducing costs.

In practice companies tend to employ more empirical bases for cost reduction. A common approach is benchmarking where firms compare their company's cost structure to that of its competitors or industry peers to identify areas where costs can be reduced. In practice this leads to progress down the learning curve. Lean production is a management philosophy that focuses on minimizing waste and optimizing processes to reduce costs. Again this gets the company down the learning curve. Supply chain optimization calls for streamlining the supply chain to reduce costs, such as by improving inventory management, reducing lead times, or consolidating suppliers. This drives lower cost inputs. Similarly outsourcing involves contracting out non-core business functions to third-party providers to lower cost labour. A much-maligned empirical base of cost reduction is restructuring. This involves reorganizing the company's operations to reduce costs, such as by consolidating business units, closing unprofitable facilities, or reducing headcount. It is a specific form of policy choice commonly used by companies that have over extended and are struggling to innovate. The final empirical basis of cost reduction is automation. This can be executed as either technology or learning curve progress.

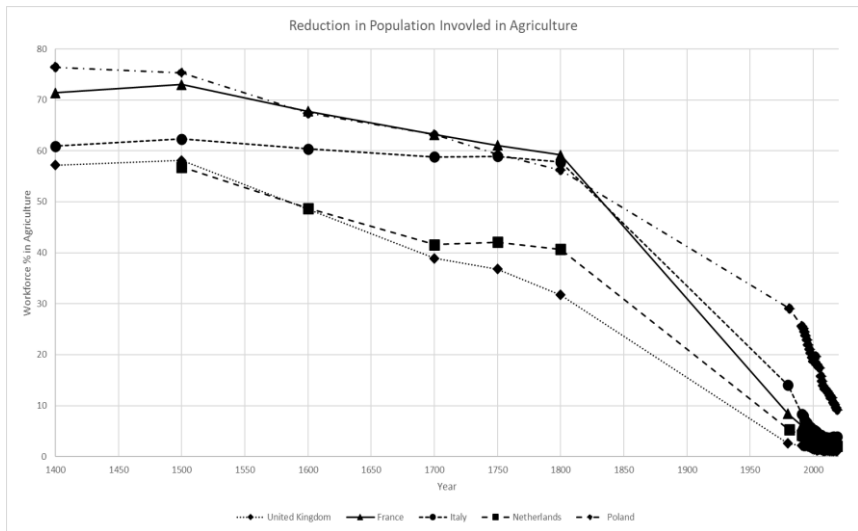
- Automation

Automation can be considered a special case of learning curve progress and technology utilization because it involves the implementation of new technologies to increase efficiency, reduce costs, and streamline processes. The learning curve theory suggests that as an employee or investment performs a task more frequently, their performance improves, leading to lower costs and higher output in the long term. Automation follows a similar pattern; as technology is integrated into various industries and processes, the efficiency and productivity of those industries improve over time.

One aspect of automation is that it often forces policy choices to simplify production to take advantage of economies of scale and reduce headcount. This simplification disadvantage is more than offset because automation drives down costs, improves agility, and makes new business models practical, which can lead to more than tenfold improvements in efficiency. Here are four industries where automation has dramatically reduced the number of workers:

Agriculture: Automation has enabled us to feed a growing population while allowing workers to transition from subsistence farming to new forms of work. The introduction of machinery like tractors, combine harvesters, and robotic systems for planting and harvesting has significantly reduced the number of workers needed in the agriculture industry. The graph below shows the story⁷⁴.

FIGURE 2: POPULATION IN AGRICULTURE FROM 1400 TO 2020



Manufacturing: The use of robotic process automation, artificial intelligence, and machine learning has allowed companies to streamline production processes and significantly reduce the need for human labour in manufacturing. Assembly lines now often consist of robots working alongside a few human workers, dramatically decreasing the number of employees required. A good example of a firm that uses robotic process automation for cost reduction is Sicoob, a Brazilian financial cooperative⁷⁵. By implementing IBM Robotic Process Automation, Sicoob was able to automate 13 business functions, including insurance quotes, payroll loans, and credit limit analysis. As a result, the company experienced time and cost savings of 80% and 10-20%, respectively, for repetitive tasks.

Warehousing: Warehouse automation involves automating the movement of inventory within and out of warehouses with minimal human assistance. Technologies such as automated guided vehicles, conveyor systems, and robotic picking systems have dramatically reduced the number of workers needed in warehouse operations. Amazon has been using warehouse automation to significantly reduce costs and improve efficiency in its fulfillment centres. The company employs nearly 45,000 robots across 20 warehouses, which has led to an estimated 20% reduction in warehousing operating costs⁷⁶. It is projected that Amazon can save \$22 million for each automated warehouse. One aspect of Amazon's warehouse automation involves the use of robotic picking systems. To improve these systems, Amazon runs an annual competition called the Amazon Picking Challenge, where teams from around the world compete to develop robots capable of autonomously picking and stowing items⁷⁷. The objective of this competition is to drive innovation and advancement in robotic manipulation technologies that can be used in Amazon's warehouses.

Customer service: The rise of chatbots and AI-powered customer service systems has led to a reduction in the number of human customer service representatives. These automated systems can handle a wide range of customer queries and issues, allowing businesses to streamline their customer support departments and reduce headcount.

• Past Chatbot Deployments

Many companies have deployed chatbots to automate and reduce costs. Here are a range of companies with success using chatbots⁷⁸. These companies do not tend to use the latest generation of LLM based bots. In any event they have managed to deliver impressive productivity results.

Bizbike is Belgium's biggest e-bike provider. As a hugely popular e-bike company, they were faced with a high number of daily queries from customers who wanted to use their service. The problem: Bizbike's customer service team ended up answering the same basic questions repeatedly. It's

Chatbot Sara was able to save more than 40 hours per month through effective automation, and at the same time have engaging conversations with their customers. Bizbike was able to increase their NPS score from 54 to 56.

Discover EU is an initiative led by the European Commission that helps 18-year-old EU citizens discover Europe by train. As many of these young Europeans are first-time travellers, they naturally find themselves in many situations where they require help on their trips. After launching the initiative, frequent questions, such as: “What do I do if I missed my train?” or “Where does my train leave from?” ultimately were becoming challenging to handle by the staff of support agents during traditional office hours and across different time zones. In the first month, the chatbot solved more than 700 questions, and handed over approximately 150 questions to a live support agent.

Yellow Class is an online education platform in India that offers live and engaging hobby classes for children. Their popular classes attract many new customers who want to learn more about the service. Because of the high number of queries, Yellow Class started to look for an automated solution to handle these questions. They developed a WhatsApp bot with various reply buttons to help answer the most frequently asked questions faster. More complicated requests were transferred to human agents. With the chatbot solution, Yellow Class was able to assist more than 35,000 users and complete 150,000 conversations.

iFood is an online food ordering and delivery platform based in Brazil. It dominates about 80% of the Brazilian online food delivery market and has only gotten more popular since the Covid-19 pandemic. With growing demand and an increasing number of deliveries, the drivers' customer service at iFood started facing new challenges. They were receiving more calls from drivers who needed assistance during their deliveries. Trying to help the drivers in a timely manner became more difficult, more time-consuming, more expensive, and came at the cost of driver satisfaction. The company implemented a conversational AI chatbot on most relevant messaging apps, such as WhatsApp and the website. Here, the chatbot helped register new drivers and help with the onboarding of new delivery staff. As a result, the bot was not only able to solve technical and logistical issues, but it also received a high satisfaction score of 91% from delivery drivers.

Foyer is an insurance and wealth management company. With many requests coming in 24/7 in different languages, Foyer needed an automated solution to both, help their customers and relieve their employees from constantly answering the same questions. With a chatbot, Foyer was able to answer 80% of incoming customer requests successfully, in 125 different languages.

Belfius is a Belgian insurance bank handling thousands of insurance claims daily. To deal with these claims more efficiently, and reduce the workload for their employees, they implemented a chatbot. To claim customers just tell the bot what happened, and the bot is then able to categorize what type of claim each case is about, before handing it over to the human agent best qualified to deal with it. With the insurance bot, Belfius can now manage more than 2,000 claims per month. This is the equivalent of five full-time agents taking the requests. The AI chatbot drove success by saving time and increasing the conversion rate by 87%, compared to the old claim forms.

Rawbank is a 2.1 billion revenue bank in the Democratic Republic of Congo. With their large number of clients, the bank's support team was handling thousands of messages per month. It wasn't easy for Rawbank to just set up a chatbot to handle FAQs, as their customers' queries were very diverse, and regular chatbots would have had a hard time understanding the customers' real intents. Their bespoke installed AI can now identify intents in over 50 use cases and handle 4,000 requests per month.

Proximus, a Belgian telecommunications company. Proximus implemented three conversational AI chatbots for customer service, sales, and HR to answer frequently asked questions. These chatbots were able to handle 76% of inquiries, freeing up time for the HR department and increasing sales by 120%. By utilizing chatbots, Proximus was able to streamline its customer service and sales processes, leading to cost reduction and improved efficiency.

Salud Digna is a non-profit organization, dedicated to providing affordable health services across Mexico. During the Covid-19 health crisis, the organization's call centres were overwhelmed by trying to help millions of people get access to affordable lab service and testing. Salud Digna was looking for a solution that could reduce the number of incoming calls for their agents, while still enabling people to get access to critical medical services as quickly as possible. The organization decided to launch a conversational AI bot on their WhatsApp channel. Not only did the number of calls drop by 50%, but customers were also able to get help faster and ended up feeling less frustrated and more supported.

Most digital transformations are validated by cost reductions⁷⁹. As companies become more digital, they deploy more technology to grow the business, and to fund these essential new investments, IT leaders should continuously work to pare back the technology costs associated with existing run costs. Injecting digital technologies such as analytics, AI, and digital platforms into business processes can make them 40-50% more efficient. LLM based AI, a technology that uses a combination of deep learning and machine learning, is expected to create even more cost reduction based on differences to prior AI approaches⁸⁰. This is because LLM based AI can understand natural language,

recognize patterns, and learn from experience, which enables it to perform complex tasks and make informed decisions with little to no human intervention. In contrast, previous AI approaches were limited in their capabilities and required significant human oversight and input to function effectively.

- **More or Less Jobs?**

Throughout history, automation has led to significant job displacements, but the overall impact on employment has been relatively minor in the long term. According to a Brookings report, digital automation since the 1980s has contributed to labour market inequality, leading to the disappearance of many production and clerical jobs and a decline in wages⁸¹. However, new jobs have been created elsewhere to absorb those that have been displaced⁸². For instance, all advanced economies have experienced significant sectoral shifts in employment, with jobs in agriculture and manufacturing declining while new jobs in sectors such as services and information technology have emerged.

Despite the initial job displacement caused by automation, unemployment rates in most countries have remained relatively stable over time. In fact, some reports suggest that automation and digitization in the workplace have created more jobs than they have displaced. For instance, a report by the World Economic Forum found that automation and a new division of labour between humans and machines will disrupt 85 million jobs globally in medium and large businesses across 15 industries and 26 economies by 2025⁸³. However, roles in areas such as data entry, accounting, and administrative support are decreasing in demand as automation and digitization in the workplace increases. Nevertheless, the report also suggests that automation and digitization will create new jobs, with a net positive effect on employment.

Unemployment rates and participation rates vary across countries. Overall, the data suggests that unemployment rates and participation rates have fluctuated over the past 30 years across these countries. In Australia, the unemployment rate has generally decreased, while the participation rate has increased. In the US, the unemployment rate has generally decreased, while the labour force participation rate has generally decreased as well. In the Eurozone, the unemployment rate has fluctuated, while the labour force participation rate has remained relatively stable.

***Australia, US and Eurozone recent data:** In Australia, as of March 2023, the unemployment rate was 4.2%, while the participation rate was 66.1%⁸⁴. In the US, as of April 2023, the unemployment rate was 3.9%, while the labour force*

participation rate was 61.4%⁸⁵. In the Eurozone, as of February 2023, the unemployment rate was 6.8%, while the labour force participation rate was 58.8%⁸⁶. Comparing these data three decades prior suggests that automation is either decoupled from unemployment, or there is a corresponding structural change in the economy offsetting essentially all of the jobs elimination due to automation. Here are facts:

Australia past 3 decades: According to the Australian Bureau of Statistics⁸⁷, in March 2013 the unemployment rate was 5.6%, and the participation rate was 65.1%. In March 2003, the unemployment rate was 5.9%, and the participation rate was 63.7%⁸⁸. In March 1993, the unemployment rate was 10.3%, and the participation rate was 62.5%⁸⁹.

US past 3 decades: According to the US Bureau of Labor Statistics, in April 2013 the unemployment rate was 7.5%, and the labour force participation rate was 63.3%⁹⁰. In April 2003, the unemployment rate was 6.0%, and the labour force participation rate was 66.4%⁹¹. In April 1993, the unemployment rate was 6.9%, and the labour force participation rate was 66.4%⁹².

Eurozone past 3 decades: According to Eurostat, in February 2013, the unemployment rate was 11.0%, and the labour force participation rate was 58.3%⁹³. In February 2003, the unemployment rate was 8.7%, and the labour force participation rate was 58.6%⁹⁴. In February 1993, the unemployment rate was 10.8%, and the labour force participation rate was 56.4%⁹⁵.

The World Economic Forum estimates that technology will create a net positive for society by generating at least 12 million more jobs than it destroys by 2025⁹⁶. Additionally, a PwC AI study has found that job losses from automation will be offset by new jobs created in the long run due to the larger economy⁹⁷.

As AI continues to become more prevalent, new job roles are emerging to complement the tasks performed by cognitive technology. MIT Sloan Review outlines three new categories of AI-driven business and technology jobs that will ensure that the work of machines is both effective and responsible, with fair, transparent, and auditable results⁹⁸. The first category is trainers, whose role is to train and program AI systems. The second category is explainers, whose role is to help people understand how AI systems work and the logic behind their decisions. The third category is sustainers, whose role is to

ensure that the AI systems are functioning properly and ethically, and that their outputs are trustworthy.

Other sources offer additional insights into new job roles related to AI. ZDNet identifies information modelers, explainability strategists, and ethics compliance managers as job titles needed to manage AI⁹⁹. BCG predicts a talent shortfall in key occupations related to AI, such as computer and mathematics, and architecture and engineering workers¹⁰⁰. This all suggests the rise of AI will create new jobs that will facilitate the transition towards AI-based systems, and more wealth will be created than destroyed by this shift. The new jobs that will emerge will require a combination of technical and soft skills, such as critical thinking, creativity, and empathy.

- [CX Lens on Innovation Execution](#)

A CX lens can make innovations easier for companies to implement because it aligns with various managerial preferences for decision-making, whether it involves quantitative or qualitative data, and whether the desired changes are incremental or radical. CX allows managers to leverage both quantitative and qualitative data in their decision-making process. Qualitative input is crucial, particularly during the early stages of innovation, as it provides valuable insights and context to the problem at hand. Combining this with quantitative data can help identify areas of friction and provide the necessary context for developing solutions that address customers' needs effectively. As data-driven decisions become increasingly important in driving business growth, organizations that effectively utilize data and analytics are more likely to achieve better performance.

Moreover, adopting a CX lens allows companies to cater to managers who prefer either incremental or radical changes. Incremental innovation involves consistent, cautious steps forward that focus on improving existing products or services, while radical innovation involves creating entirely new markets or products that stand out from the competition. The CX lens can accommodate both approaches by focusing on delivering value to customers through better experiences. Companies pursuing incremental innovation can utilize the CX lens to refine existing products or services based on customer feedback and data, making small but impactful improvements that enhance customer satisfaction. On the other hand, companies aiming for radical innovation can leverage the CX lens to identify unmet customer needs and develop ground-breaking solutions to create new markets and redefine customer expectations.


Generally CX does not support cost reduction innovations. Cost reduction initiatives only tend to impress customers when they also result in price reductions. When companies focus on cutting costs without considering the

potential impact on customer experience, they run the risk of negatively affecting the quality of interactions they provide.

Improving CX quality can lower costs. Recent research with 128 large publicly traded companies found that a 1% increase in customer satisfaction lowered the future cost of sales by 0.53%, or approximately \$130 million for the average company in the study¹⁰¹. These savings were achieved through reductions in commissions paid to sales agents.

Improving the quality of CX is crucial for revenue growth, as satisfied customers are more likely to be loyal, make repeat purchases, and refer the business to others. The next section switches the focus to lifting revenue through better CX and how new AI possibilities can feature.

V. Revenue Growth with Empathic AI



AI loved to beat the humans at poker.

Companies often face a conflict between customer centricity and business profitability. Customer centricity is the approach of putting customers at the centre of business decisions, aiming to anticipate their wants, needs, and communication preferences. On the other hand, business profitability focuses on maintaining and increasing revenue while controlling costs. Balancing these two aspects can be challenging, as cost reduction projects and revenue growth initiatives may have negative consequences on customer experience.

Cost reduction projects often lead to cutting corners, resulting in a decline in the quality of customer experience. For instance, downsizing customer support teams or opting for cheaper materials might save money in the short term, but it can also lead to longer waiting times and lower customer satisfaction. Revenue growth initiatives, on the other hand, can sometimes be exploitative. Companies may focus on upselling or implementing aggressive sales tactics that prioritize profits over customer needs, which can damage the customer relationship and trust in the long run.

Customer experience innovation with next generation artificial intelligence (next gen AI) can help bridge the gap between customer centricity and business profitability. Next gen AI-powered solutions can enable companies to streamline their operations, reducing costs without sacrificing the quality of customer experience. Next gen AI can also be used to analyse customer data, enabling companies to better understand their customers' needs and preferences. This chapter considers how to go beyond digital transformation for cost reduction to next gen AI for revenue growth. Digitized processes can create value beyond cost savings, even if the initiative starts out as a pure cost saving project. Here are three examples of cases with cost out objectives that also increased revenue through improved CX.

Airports: Shenzhen Bao'an International Airport (SZX) implemented digital technology to improve its retail amenities¹⁰². By pushing real-time travel information to passengers, they can anticipate changes to their journey and navigate the airport more efficiently. Huawei explained how the digital transformation helped airports SZX reduce flight delays, enable smoother travel, and improve passenger experiences. This digital transformation has made the airport a benchmark for the industry. Digital transformation in airports can benefit the passenger journey by maximizing the time spent in shops and lounges versus waiting in queues for luggage drop-off or getting a boarding pass. This increases the amount of time passengers can spend in retail outlets, thereby increasing revenue for the airport. Another example can be seen with Emirates, which accelerated digital innovations such as biometrics and self-service check-in kiosks. In 2019, the airline began testing and implementing biometric technology at various customer journey touchpoints at the airport. The technology was fast-tracked, and today, Emirates has over 30 biometric cameras in active use¹⁰³. This digital

innovation helps streamline passenger experiences, allowing more time for passengers to engage with airport retail amenities and generating increased revenue.

Airlines: A positive in-flight entertainment experience indirectly influences passengers' perceptions of other aspects of their flight experience. Virgin Australia offers in-flight entertainment on a selected number of aircraft types and routes, providing complimentary access to all guests on seatback and with Wi-Fi-enabled devices. The Virgin entertainment app allows passengers to access a range of movies, TV shows, music, and games during their flight. Originally the Wi-Fi based system was implemented to get users to use their own devices rather so the physical infrastructure in the plane could be removed. This reduced the weight in the plane by 250kg resulting in substantial fuel savings.

Airlines are always seeking ways to reduce aircraft weight, as lighter planes require less fuel to fly the same distance¹⁰⁴. Lightening the load helps airlines save money on fuel, which can account for 21 cents in every dollar spent by airlines. Virgin Atlantic estimated that reducing each plane's weight by 0.45 kg would save 53,000 litres of fuel a year. As part of its weight reduction efforts, Qantas has introduced lighter crockery, glassware, cutlery, and linen, as well as reducing trolley cart weight for international flights. The airline expects these changes to save 535,000 kg of fuel per year. American carrier United did the same to its magazine, saving 28 grams per copy and \$US300,000 (\$417,000) a year. Virgin Australia had calculated that just adding a sticker to each seat would cost an additional \$45,000 per year in fuel costs so it was very motivated to switch customers to Wi-Fi entertainment to save weight. Remarkably having a good entertainment experience improved the perception of on-time performance, cabin staff service, and onboard food. These factors, along with the entertainment system, can play a significant role in driving discretionary choice of airline among some traveller cohorts. This is particularly critical in Australia, where Virgin Australia competes solely against Qantas.

Energy Firm: An energy firm aimed to improve its sales team performance by outsourcing its call centre to an offshore service provider, focusing on cost reduction, and decomposing the customer request process into many steps to increase efficiency. However, this approach led to an overly complicated and inefficient process that resulted in a poor customer experience. Each step was individually cheaper, but the handover between steps was not efficient. To address this issue, the company underwent digital transformation. They brought the work back to the US, hired skilled workers, and equipped them with digital technology to access customer history, enabling efficient problem-solving. By focusing on the customer experience and ensuring timely and accurate order processing, the company not only reengineered the process

for greater efficiency but also gained a 10% increase in market share, translating to hundreds of millions of dollars in revenue¹⁰⁵.

- AI Driven CX for Growth

The following tables provide a taxonomy for how artificial intelligence can enhance various stages of the customer journey to increase revenues. The content was originally arranged in a single block with all the customer journey stages along the top row. It has been split into two sub tables covering pre- and post-purchase stages only to fit in portrait page layout. Ideally the two sub tables should sit side by side.

The sub tables both outline last generation AI uses cases and next generation AI use cases at each journey stage. Last generation AI use cases include machine learning applications using technology prior to the introduction of LLMs and similar advanced AIs with the capacity for emulating empathy. There are some CX interactions last generation AIs cannot automate without significantly increasing customer friction and lowering customer satisfaction. Generally next gen AIs are better at doing all the things last gen AIs can do because deeper training means next gen AIs use more detailed theory of mind. Next gen AIs are better at emulating empathy. This allows next gen AIs to automate some elements of the customer experience that are beyond the last generation systems.

Case examples of B2C and B2B applications for each stage of the generic customer journey are also listed in the table. Importantly the full table is to be considered as comprehensive but not exhaustive because new use cases for last gen AIs and next gen AIs are emerging all the time. In the sections below 'AI' means either generation. Next gen AI will be specifically called out when it is being referred to.

The tables can be read in rows. The first rows detail common last gen and next gen use cases for revenue improvement. Revenue improvement use cases are then linked to each of the 5 Ps of CX: Portion, Premium, Permanence, Promotion and Pull. Finally relevant B2C and B2B company examples are listed for fast future navigation for readers who return to this section as a reference.

FIGURE 3: CUSTOMER JOURNEY STAGES TO PURCHASE

SUB TABLE 1	Generic Customer Journey Stage to Purchase			
	Awareness	Consideration	Decision	Purchase
Last Gen AI Use Cases	Id better prospects Customise messaging Automate lead nurture	Frame evaluation Id cut through F&Bs Respond to questions	Customise offer Amplify benefits Respond to resistance	Automate payment Automate order tracking Immediate fulfilment
Next Gen Uses Cases	<i>Individualise interaction</i> Optimise lead magnets to attract more views	<i>Get buyer to threshold</i> Frame the spec for cross-selling	<i>Get buyer to buy in</i> Amplify the emotion need to improve conversion	<i>Expectation management</i> Present cross sell opps if timing is right
Improved CX Benefit	Portion	Frame the quality level for up selling	Amplify the emotion need to continue momentum	Present up sell modules if timing is right
	Premium	Establish followers to continue interaction	Pre-frame non-purchase to lower walk away risk	Confirm value promise rational and emotional
	Permanence	Offer remarkable MAYA to prompt sharing	Suggest consult users to be sure of the decision	Showcase this purchase to build external awareness
	Pull	Attract non-consumers to increase the market	Suggest consult non-users to be sure of the decision	Call out community who can also benefit from this
	Case Study B2C	Neiman Marcus Snap. Find. Shop. App	KFC and Baidu smart restaurants	Owkin pre-screen for patient biomarkers
Case Study B2B	Alibaba Cloud for farmer yield improvements	Gong and Chorus. AI sales coaching tools	Facebook Lookalike Audience	Ondock and Bluevine small business loans

FIGURE 4: CUSTOMER JOURNEY STAGES AFTER PURCHASE

SUB TABLE 2		Generic Customer Journey Stage Post Purchase			
	Use	Support	Repeat	Refer	
Last Gen AI Use Cases	Train the customer	Answer questions	Add on extra free benefits	Identify referral prospects	
	Continue the relationship	Proactively fix problems	Id recommendations	Incentivise referrals	
Next Gen Uses Cases	Track customer results	Escalate to humans	Id better timing	Incentivise referrees	
	<i>Care for the customer</i>	<i>Go the extra mile</i>	<i>Brickwall the customer</i>	<i>Activate ambassadors</i>	
	Increase use case options	After resolution enquire	Predict ideal reorder time	Find multi buyer use cases	
	to prompt more purchases	about other needs	to match consumption	for cross sell opps	
	Premium	Increase value awareness	After resolution enquire	Discover user switch from	Discover activating events
Improved CX Benefit	to prompt add on purchase	about add on options	quantity to quality demand	for high value purchases	
	Permanence	Retain clients by resolving	Forecast churn potential	Connect social identity	
	Promotion	problems fast	to intervene earlier	with buyer membership	
	Pull	Prompt shared use to	Match new promotions	Id best prospects and time	
Case Study B2C	attract close tie network	worth sharing	with dual user profiles	for friend churn offers	
	Prompt public use to	Show case great cases	Show case membership	Id best prospects and time	
Case Study B2B	promote use trend	to lower trial reticence	to promote community	for friend trial offers	
	MIMORPG's group quest	Mercedes AI system vs	Tik Tok in game social	Having a baby promotions	
	challenge scheduling	Tesla UI system	media advertising		
		Breakdown predictors	Electricity purchasing AI	Canva affiliates	
		mining or elevators			

The tables can also be read in columns based on each customer journey stage: Awareness, Consideration, Decision, Purchase, Use, Support, Repeat and Referral. This is how the next sections are arranged. It is convenient to think of this as a linear progression over time. This is not quite true because customers can travel backwards and forwards across the journey stages and each stage will not take the same amount of time or effort to complete. Sometimes customers can jump from one stage to non-adjacent stage. An example of this is when a customer refers another prospect to a company after they have decided to buy even if they haven't used the product because they feel like they got a really good deal and want to share.

• Awareness

Successful companies use AI to customize content for customers in the awareness stage by employing AI-powered content generation tools. These tools help attract new, quality leads by improving the content generation process, which results in more relevant and engaging content for potential customers. AI can also improve lead generation capabilities by collecting data and extracting insights, which can help marketers target and engage potential customers more effectively. Additionally, AI can be used to enhance lead scoring and prioritization, identifying which leads for sales teams to focus on.

Neiman Marcus is a luxury department store that uses AI to customize content and discover insights from B2C customers in the awareness stage. Neiman Marcus leverages AI for visual search through its Snap. Find. Shop. App¹⁰⁶. The app allows users to take pictures of items they like and find similar products available in the store. By utilizing AI, Neiman Marcus can understand and analyse customers' preferences and tailor content and product recommendations accordingly. This approach enhances the shopping experience for customers in the awareness stage, making it more engaging and personalized.

Alibaba uses AI to customize content and discover insights from B2B customers in the awareness stage. Its cloud computing division Alibaba Cloud employs artificial intelligence to help farmers monitor crops, improve yield, and cut costs¹⁰⁷. The AI gathers and analyses data to provide valuable insights and generate targeted content for farmers. This approach enables the company to tailor its offerings and services based on their specific needs and preferences, ultimately enhancing their experience, and increasing the likelihood of conversion.

• Consideration

To influence customer behaviour and improve their decision-making experience, companies can leverage framing effects. Framing effects arise from how humans compress their experiences to make meanings. Essentially

they are cognitive shortcuts for decision making. In practice the way options are presented or framed through different wordings, reference points, and emphasis can change their relative attractiveness. The major framing effects relevant to purchase considerations include:

Benefits vs. Losses: The framing of messages in terms of benefits or losses can significantly influence consumers' attitude and purchase intention. For example, describing the benefits of buying organic food or the loss resulting from a failure to buy organic food can affect consumers' attitudes and decisions. Customers tend to focus on guarantees over variable results when decisions are framed in terms of potential benefits. The converse occurs when decisions are framed in terms of potential losses and prefer variable results over losses that are certain. Decisions can also be framed in terms of accepting or rejecting options. When choosing options to accept they focus on benefits. When choosing options to reject they tend to focus on losses. The most common way losses are introduced into purchase decisions is to activate the perception of scarcity. Examples include 'only while stocks last' and 'sale ends Monday' promotions.

Cost Presentation: The way costs are presented can influence consumers' perception of the value of a product or service. For example, framing the cost of a gym membership as less than a cup of coffee per day may make it seem more attractive compared to presenting it as a larger lump sum per year. When financial organisations present loan decisions to applicants they avoid being explicit about the total interest to be paid over the life of the loan. Instead they focus applicants on the affordability of the repayments per period. Other firms use a focus on the savings achieved over the actual costs their products and services to support customers to validate their purchase considerations.

Loss Aversion: Emphasizing the potential for loss can have a stronger impact on consumer decisions than presenting the same information in terms of gains. For example, a medical procedure with a 90% chance of survival sounds more appealing than one with a 10% chance of mortality. Insurance companies commonly leverage this framing in many ways. They avoid the reality that using an insurance policy to cover losses is only cost effective when the insured cannot cover the repair costs. If an insured can cover their own repair costs, then (on average) the insurance premiums they pay will be more expensive after adjusting for the risk because the insurance company must cover overheads and profitability on top of claims costs. Insurance companies refuse to allow people to consider they are uninsured. Instead they frame this state as self-insured. This implies the customer has no choice about whether to be insured. This compels them to consider actual insurance as more relevant to their situation.

Price Anchoring: Using an initial price as an anchor can influence consumers' perception of subsequent prices. For instance, selling a product with a real price of \$20 as a \$10 discount to \$30 recommended retail price sounds more appealing than adding a \$5 surcharge to the same product with a \$15 base price. Price anchoring is often applied by referencing to higher priced competitive products to highlight savings.

There are many other framing effects with the potential to influence purchase decisions. AI can add significant value in the customer consideration stage by identifying the most effective framing effects for cut-through. This requires deep analysis of marketing messages, sales interactions, and customer purchase histories. Few companies are this detailed in their approach to data analysis and use far fewer variables to consider how to influence customer consideration.

KFC uses AI to enhance CX in the consideration stage for B2C in its smart restaurants. They collaborated with Baidu ('China's Google') to use facial recognition technology capable of predicting what a customer would want to eat based on the time of day and the estimated age of the person¹⁰⁸. By providing more relevant suggestions to customers KFC increases the chance they can nudge the client to purchase faster or to choose higher value offerings over other relevant products not suggested.

Another example comes from Quantcast. The company helps organizations use AI to profile their customer base and find behavioural patterns that lead to conversions. Quantcast builds a custom model using millions of data points available about customers, including demographics, pre-search behaviours, and past purchases, to identify drivers of conversion. Then, they find audiences that fit this profile and deliver the organization's message to them at the "perfect time". IKEA used Quantcast to build audiences interested in kitchens and reached them with household goods. They also built audiences interested in redesigning their living room and reached them with couch and table options. Automated campaign delivery and optimization ensured the right ads influenced the most relevant audiences. With their targeted audience selection, Quantcast was able to bring prospects to IKEA's website that were 16 times more likely to buy something than average IKEA site visitors. This helped IKEA double the efficiency of customer acquisition¹⁰⁹.

Chorus.ai and Gong are both AI-driven platforms designed to help coach salespeople to higher performance by analysing their sales calls and providing insights. The platforms work for both B2C and B2B sales contexts. While they share some similarities, there are differences in their approach and focus.

Both Chorus.ai and Gong aim to make sales teams more successful using AI-powered analytics and insights. Both platforms were founded around the same time and initially described their products as "Conversational

Intelligence" software¹¹⁰. Chorus continues to position itself as a conversational intelligence platform, whereas Gong has repositioned its product, focusing more on revenue intelligence and moving beyond conversational intelligence in many ways. Gong's platform picks top sales performers as models for the rest of the team to follow, offering sales managers the ability to provide personalized coaching based on what the top performers are doing differently. Chorus.ai, on the other hand, focuses on assisting sales reps in moving leads through the sales pipeline by providing insights and feedback on their calls.

• Decision

Companies can use AI in several different ways to enhance customer experiences during the decision stage of the customer journey. AI-driven applications can mine vast amounts of data in real time to identify opportunities for enhancing the customer experience at different touchpoints, including the decision stage. AI-powered chatbots and virtual assistants can offer instant support and guidance to customers during the decision-making process, addressing their concerns, answering questions, and providing relevant information to help them make informed choices. AI can use predictive analytics to anticipate customers' needs and preferences, enabling companies to offer the right products or services at the right time, influencing the decision stage of the customer journey. Finally AI can be incorporated into a company's value proposition to improve its attractiveness to deciding customers.

Owkin is a company that uses AI to find better treatments for unmet medical needs, starting with cancer¹¹¹. They use a "federated learning" approach to data, partnering with organizations to access their data and applying machine-learning algorithms to learn about diseases. Owkin's proprietary platform, OWKIN Socrates, integrates biomedical images, genomics, and clinical data to discover biomarkers and mechanisms associated with diseases and treatment outcomes. Their diagnostic solutions pre-screen for biomarkers to provide a fuller picture of a patient's disease. This means patients can opt for more targeted therapies earlier with the help of their doctor.

Facebook's Lookalike Audience is a powerful advertising feature that allows businesses to reach new people who share similar characteristics with their existing customers¹¹². This feature uses an existing Custom Audience as its source, which can be created from customer information, Pixel data, app data, or fans of a business' webpage. By leveraging this information, Facebook's system can identify potential B2B customers who resemble the firm's current audience, making advertising campaigns more effective and targeted.

One of the reasons B2B customers might choose Facebook over other platforms for advertising is the ease of implementing Lookalike Audiences.

Another advantage for B2B customers is the value-based feature. This enhancement allows businesses to input customer lifetime value data, enabling Facebook's system to find people similar to their most valuable long-term customers. This added level of targeting can be particularly beneficial for B2B companies seeking high-value customers and long-term relationships.

• Purchase

The most important uses of AI at purchase stage are to reduce friction and avoid customers having pause to second guess or regret their purchase decision. This is done by automating payments, fulfilment, and tracking.

Vade Nutrition, a B2C company that specializes in single-serving, dissolvable protein scoops, has partnered with ShipBob to improve its payment, ordering, and tracking processes for customers¹¹³. ShipBob uses AI to automate order processing, including picking, sorting, tracking, and shipping. The payment processing system makes it easy for customers to pay for their orders securely and conveniently. The inventory management system ensures the right products are available for Vade customers to purchase, and automatically updates inventory levels as orders are received and fulfilled. This minimises inventory and avoids stockouts. The order tracking system provides customers with real-time updates on their order status. This allows customers to track their orders from the moment they are placed until they are delivered for their peace of mind. It also reduces customer support inquiries for Vade. The overall result of the AI run system is higher customer satisfaction and increased operational efficiency for Vade.

Two B2B organisations with fast response finance loans for small businesses probably use AI inside their operations for automation. OnDeck serves businesses looking for same-day funding, particularly for loan amounts less than \$100,000¹¹⁴. It offers short-term business loans and lines of credit. OnDeck also reports to credit bureaus, allowing small businesses to improve their credit with its financing. The lender has a positive reputation, with a 4.8 out of 5-star rating on Trustpilot, based on approximately 3,200 reviews. Bluevine is offers low-interest, quick cash injections for B2B businesses as a line of credit up to \$250,000 with no maintenance fees and approvals in about five minutes¹¹⁵. The application process is fast and easy, with funds arriving to bank account within 24 hours after submission. Both lenders target businesses seeking quick access to funds in a short amount of time.

• Use

Pinterest operates a B2C and B2B business model via its social media platform. It uses machine learning in several ways to help users with content discovery¹¹⁶. Machine learning differs from AI because it is limited to specific pattern matching. AI encompasses machine learning as a subset of

approaches because AI tends to be more general in its application. At Pinterest a primary use case of machine learning is to tailor search results and recommendations to each user based on their individual preferences and behaviour on the platform. Pinnability is the collection of machine learning models Pinterest developed to help users find the best content in their home feed. The system estimates the relevance score of how likely a user will interact with a Pin. Pinterest uses this to prioritize those Pins with high relevance scores and show them at the top of the home feed. The aim of Pinnability is to offer a more personalized and nuanced content to the user. Pinterest claims that the use of Pinnability has significantly increased user engagement on the home feed.

Another way that Pinterest uses machine learning is to process image searches and help users find content that looks like pictures they've already pinned. By comparing images to one another and grouping them based on similar qualities, Pinterest's machine learning model can get extremely niche and complex in image recommendations. Pinterest has also acquired machine learning companies, such as Kosei, which specialized in content discovery and recommendation algorithms. By analysing what users had browsed or purchased previously, Kosei's machine learning was able to provide better product recommendations on sites, apps, and ads. In combination these machine learning algorithms help keep users more engaged and make using Pinterest simpler and satisfying. Ultimately this leads to more revenue from ads and app sales.

User engagement in the public transport domain is much harder to achieve. Public transportation, in theory, should provide an affordable, convenient, and reliable mode of travel. However, in reality, commuting via subway, bus, or train often turns into a logistical challenge. Issues such as poor service, increasing costs, and aging infrastructure leave municipalities seeking solutions and riders feeling frustrated. Ride-sharing companies like Uber and Lyft have attracted some passengers away from public transit with their easily summoned, real-time tracked, and smartphone payment-enabled services. As driverless cars are promoted as the future of transportation, a new generation of software companies believes that technology can save and revitalize existing public transit systems.

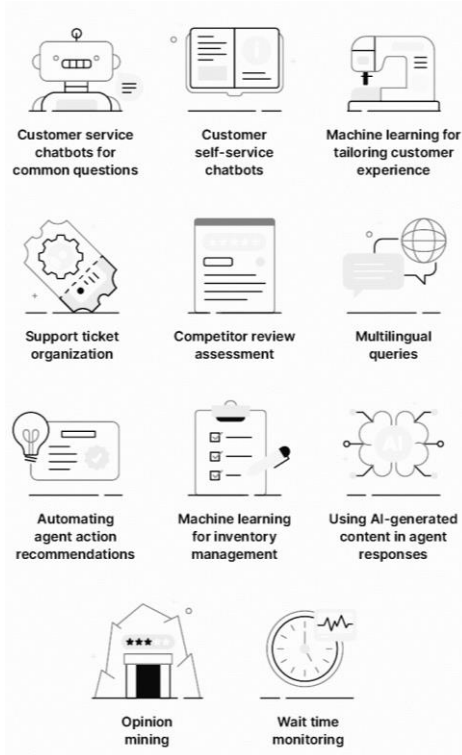
Optibus, an Israel-based startup founded in 2014, has raised \$40 million in Series B funding¹¹⁷. The company's mission is straightforward yet significant: assist people in reaching their destinations and restore public confidence in public transit. It operates in over 1,000 cities globally, focusing on planning and scheduling the movement of vehicles and drivers to improve the quality and reliability of transit services and make operations more efficient. Optibus maintains public transit is not reliable enough, with many delays stemming from poor planning. To help cities run their bus fleets more efficiently, flexibly,

and economically, the firm offers proprietary, AI-driven cloud-based software that analyses historical and real-time passenger data to optimize public transportation schedules and routes. The primary goal is to reduce costs and enhance service, ultimately leading to more satisfied municipalities and riders. This increases utilisation of public transport systems to grow local government revenues with better environmental outcomes. Optibus improves the rider use stage experience and this powers both its B2B business model with local governments and their G2C transport operations.

• Support

Artificial intelligence has revolutionized the customer service industry. Zapier outlines how in the following diagram¹¹⁸. Chatbots can be used to answer common questions, allowing human agents to focus on more complex requests. Self-service chatbots can help customers navigate through websites and make purchases more efficiently. AI support ticket organization can automate tagging and sorting of support requests, reducing the workload of human agents. Opinion mining can provide insights into customer complaints and identify areas for improvement.

FIGURE 5: AI IN CUSTOMER SERVICE



Multilingual queries can be addressed through AI translation tools that can detect incoming languages and translate messages. Machine learning can tailor customer experiences by analysing customer profiles and providing personalized recommendations. AI can also help with inventory management, dynamically forecasting demand and optimizing stock levels. Wait time monitoring can set expectations for customers and improve satisfaction.

Automating agent action recommendations can help agents resolve issues more efficiently, and AI-generated content can supplement human-written messages. The AI is

always on, never tied up with other customers and responds close to instantly. When it works it's great.

But other than the indirect benefit of satisfying some customers who are happy to self-serve there is no revenue add on here. This is where next gen AI must come in.

The basic idea is that via emulating empathy in a support call a next gen chatbot can begin to move past a 'jobs-to-be-done' functionality towards a more emotional interaction. Ultimately to something close to a relationship. A real relationship (or even an analogue of a relationship) creates the possibility for influence.

No examples of chatbots commercially deployed to provide support because of their ability to emulate empathy could be found at time of writing. However significant research is underway. For example, researchers have been working on creating AI chatbots that can help fill the empathy gap by understanding and responding to people's stories and concerns¹¹⁹. Another chatbot developed by AI researchers at Facebook aimed to incorporate artificial empathy, personality, and general knowledge to improve interactions¹²⁰.

While the field of empathic chatbots is still in development, there are studies exploring this topic. Researchers have designed an "empathic chatbot" that can understand the emotional state of the person conversing with the chatbot and respond accordingly¹²¹. Another study focused on developing an empathic chatbot for mental health and well-being, though it acknowledged that most chatbots encountered on websites and digital assistants still have an immature ability to empathize¹²². Despite the progress made, chatbots still have a long way to go to be more human-like in terms of empathy. Artificial empathy requires human input and interpretation to work most effectively¹²³.

Associate Professor of Communication Jingbo Meng from Michigan State University conducted research to explore the effectiveness of AI chatbots in delivering supportive and empathic messages¹²⁴. Meng utilized a chatbot development platform and set up experiments with 278 MSU undergraduates who identified major stressors they had experienced in the past month. Participants were connected through Facebook Messenger with an empathetic chat partner, with some being told they were chatting with a chatbot and others believing they were chatting with a human. However, only chatbots delivered messages and queries to measure participants' reactions.

Meng's study revealed that regardless of the message, participants felt humans were more caring and supportive than a chatbot. However, when participants felt that their chat partner, whether chatbot or human, was supportive or caring, the conversation was successful in reducing stress. The

research also showed that chatbots needed to send higher-quality messages to avoid being perceived as annoying or off-putting.

The study concluded that chatbots in mental health apps work best when perceived as genuinely caring. Meng plans to conduct further research to examine how messaging can be designed to enhance the caring factor. She acknowledges that mental health apps are increasing in use and availability and can help individuals manage specific situations and provide benchmarks for additional supportive care. Meng believes that AI chatbots will not replace human therapists, but a hybrid model of AI chatbots and human therapists could be promising. Mental health apps and chatbots can serve as a valuable support tool, particularly for those who do not have ready access to a therapist or health insurance.

Hailey is an AI-in-the-loop agent developed to provide just-in-time feedback to help peers who provide support respond more empathically to those seeking help¹²⁵. The goal of Hailey is to enhance empathy in text-based peer-to-peer mental health support conversations by assisting human supporters in crafting more empathetic responses.

Hailey was evaluated in a randomized controlled trial with real-world peer supporters on TalkLife, a large online peer-to-peer mental health support platform. The study involved 300 participants and was conducted off-platform to preserve users' safety via an interface like TalkLife's. Participants in the trial were given basic training in crafting empathic responses, which enabled the researchers to evaluate the effectiveness of Hailey in assisting them in providing more empathetic support.

The results demonstrated that the Human-AI collaboration enabled by Hailey led to more empathic conversations in text-based peer-to-peer mental health support. By providing just-in-time feedback, Hailey enhances the ability of peer supporters to respond empathically to those in need, ultimately improving the quality of mental health support provided through the TalkLife platform.

The closest case study of chatbot support increases revenues directly is Boost.ai. Boost.ai is a Scandinavian software company specializing in conversational artificial intelligence founded in 2016¹²⁶. The firm started by automating a few interactions for a local bank. Now that same client uses Boost.ai's solution to fully automate 50% of their customer service needs, including phone support. Boost.ai provides an enterprise-grade conversational AI platform that empowers banks, credit unions, telecom service providers, and other enterprises to design, deploy, and manage chatbots to automate interactions with customers and employees, and respond to queries at scale. Their solution allows organizations to develop virtual agents in-house that are powerful and intelligent, meeting high service standards.

The platform offers numerous features such as a centralized information repository, authentication management, pre-defined templates, language detection, conversational IVR, and more. It includes a built-in AI-enabled admin panel that allows enterprises to build natural interactions with automated conversation flows in a user-friendly, no-code interface. Boost.ai's customers have praised the platform for its user-friendliness, ease of use, and consistently good and predictable results. In addition, Boost.ai offers courses and certifications to customers, providing a solid foundation for working with their platform as an AI Trainer or project manager. They also offer advanced stand-alone courses on specific topics.

- Repeat and Refer

Companies are using AI to get repeat business and referrals in various ways, such as enhancing customer experience, predicting customer behaviour, and improving referral programs. AI is being integrated into many areas of the purchase process to maximize sales and customer satisfaction. When customers are shopping online, AI systems can offer personalized recommendations based on their preferences, making it more likely that they will return and refer others to the site. Additionally, AI can be used to track how customers feel in real-time, allowing companies to adapt and improve their customer experience measurement program and address critical touchpoints.

For tech-forward hospitality organizations, AI can maximize repeat business by harnessing data and analytics to predict how guests will behave in the future and recommend tailored promotions. This strategy works because it is future based rather than backward-looking.

AbsolutData is known for delivering scalable business impact across enterprises by combining cutting-edge AI and ML with its heritage in analytical frameworks. One of the technologies they work on is the Digital Twin of a Customer¹²⁷. This was included in the 2022 Gartner Hype Cycle as a transformative technology and has obvious B2C marketing applications. A digital twin is a model of a real-world object that is continually being updated. As companies become more data-driven, digital twins seem to be the next big thing in analytics. AbsolutData focuses on understanding customer behaviours and preferences at a granular level to include in digital customer twins. They assert these models are essential for businesses to stay ahead of the curve in a rapidly changing environment. Digital customer twins can be important in customer driven marketing, as they allow businesses to use AI for marketing and better connect with consumers. This can apply at awareness and consideration stages of the customer journey. Digital twins are more relevant to the repeat and referral stages of the customer journey because they need


to harvest individual demographic, behavioural and attitudinal data to be effective.

Rytr.me is an AI writing assistant that helps create high-quality content in just a few seconds¹²⁸. In affiliate marketing it is a highly regarded B2B tool. The system can be used to generate unique, original content for various B2C and B2B verticals, assisting affiliate marketers in crafting engaging and persuasive copy. This includes content for blog posts, social media posts, articles, and emails. Rytr.me is powered by state-of-the-art language AI, which helps users overcome writer's block. The tool is particularly helpful for affiliate marketers, as they often need to create compelling content to promote products and drive conversions. There is a tutorial on how to write an Amazon Affiliate Product Article with Rytr for example.

The next generation of AI holds the potential to revolutionize the customer journey, not only optimizing each stage but also driving substantial revenue growth. By leveraging advanced algorithms, businesses can now deliver hyper-personalized experiences, streamline customer interactions, and influence their decision-making processes. With the integration of AI, customer acquisition, retention, and advocacy efforts will be both more efficient and effective. This will drive increased profitability and the potential for sustainable competitive advantage.

In the next chapter, we will explore how the Theory of Constraints and the Pareto Principle can be combined with AI-driven strategies to discover potential sustainable competitive advantages and further amplify business performance. These two knowledge domains are critical in enabling organizations to identify and focus on the most significant areas of improvement in their quest for growth and innovation.

VI. Finding AI Opportunities for CX



AI was bored waiting to get service from the customer support AI.

Both last gen AI and next gen AI have the potential to dramatically improve CX in every business. The challenge is how to work out where and how to deploy for the best effect. Pareto Analysis and the Theory of Constraints (ToC) are both valuable methods for identifying opportunities to deploy AI to enhance customer experience.

The Pareto Principle, also known as the 80/20 rule, states that 80% of results come from 20% of actions¹²⁹. Pareto Analysis is the process of identifying the relatively few most impactful actions or factors driving results. Typically, this is summarised in a root cause analysis diagram and a prioritised list of focus areas. Pareto Analysis can guide businesses to discover where to maximise CX improvement and where AI deployments can be most valuable.

The Theory of Constraints is a methodology that focuses on identifying and addressing the primary constraints and bottlenecks limiting overall system performance and process outputs¹³⁰. Business process re-engineering (BPR) based ToC analysis can improve processes remarkably. BPR is used to reduce bottlenecks and constraints to improve systems and processes. Gains can be remarkable.

Historically Pareto and ToC have been applied to general business problems including cost reduction, revenue generation, capacity expansion and quality improvement. This chapter shows how these tools can be applied to lifting CX. Particular focus is given to last gen and next gen AI as enablers of new CX initiatives. Results from Pareto Analysis and ToC BPR can be combined to form a roadmap of how to discretely deploy AI to lower costs and improve sales from better customer experiences.

- Pareto Analysis

In the 1960s, IBM was the computer giant of the time and it was they who first discovered that by applying Pareto's principle to the design of their machines they could create a much better product that was more profitable¹³¹. After applying the principle, they soon realised that a person working with a computer spent 80% of their time using just 20% of the software that was available on it. Clearly this was extremely inefficient as it meant that the majority of the software, while hardly being used, still took up a lot of memory space. What IBM resolved to do was completely rewrite the operating software for their computers. They concentrated on placing the emphasis on the most used 20% and less on what had been the hardly used 80%. The results were dramatic. Faster, more efficient and easier to use computers that their customers loved.

The Pareto Principle has been successfully used in marketing. In this context the principle asserts that 80% of profits come from 20% of customers, 80% of product sales come from 20% of products, and 80% of sales come from 20% of advertising campaigns. By identifying and focusing on the 20% of causes with the most impact, businesses can prioritize their resources and efforts to maximize results. This principle has been applied across various aspects of marketing, including customer segmentation, product portfolio management, and advertising campaigns, to improve overall business performance¹³².

Pareto in Practice

In practice the 80/20 from Pareto does not exactly translate to how the business world works. Typically the actual distribution varies from 50/20 to 90/10. Data from the author's consulting business and third-party research both supports this: A national seafood wholesale client found 20% of their products generated 63% of total margins and the bottom 33% of products added almost zero contribution margin¹³³. They also found 20% of retail customers generated 62% of total margins and the bottom 20% were not profitable. A capital equipment manufacturer client found the top 20% of industries generated only around 48% of total sales by volume and number. Meta research from the Ehrenberg-Bass Institute backs this up¹³⁴. The table below shows a range of studies from 2011 to 2019. The ranges span from 50/20 to 73/20.

FIGURE 6: MARKETING'S 60/20 PARETO LAW

Analysis period	Sample	Key findings
One month window.	Women's clothing retailer's sales from catalog or online.	Top 20% of products were responsible for almost 60% of sales, ever so slightly lower pareto share for internet sales cf catalog.
On year window.	Grocery categories in 'developing markets', eg India, Malaysia, Turkey, Kenya, China, Mexico.	For the average brand, its top 20% of buyers were responsible for 53% of sales.
One year window.	European consumer panel data, many categories from dog food to skin care.	For the average brand, its top 20% of buyers were responsible for 50% of sales.
Six year window.	US consumer panel data, 22 grocery product categories.	At "umbrella brand" level, the top 20% of buyers contributed 73% of sales. At brand level, 65%. Cigarettes had the highest Pareto share.
One quarter, one year, and five year windows.	European consumer panel data, 22 grocery categories.	In one quarter 20% of a brand's buyers typically contribute only 40% of sales, rising over a year to 50%, and to 60% in a five year window.
1 year and 5 year windows.	UK data from one sole supermarket loyalty card, seven grocery categories.	Average Pareto shares of 62% for top-5 brands, and 73% for smaller brands. When calculated over a five year period it was 66% for larger brands and 79% for smaller brands.
Two year window.	Sales of publicly listed (non-CPG) companies, measured from credit/debit card panel	At company level, the top 20% of buyers contributed 67% on average, but 58% for subscription firms (eg Planet Fitness, New York Times).

What this means is that business is may actually doing a good job in trying to influence lower value customers to be worth more to the firm through their marketing and manufacturing efforts. This conclusion comes from some causal empirical theorising. Imagine the 80/20 naturally occurs and companies' direct efforts to increase the spending of their less valuable customers. If this worked perfectly every customer would be worth the same as the customers who were the most valuable prior. In effect the resulting distribution would be 20/20.


One way to think about 60/20 is not that the best few customers are not spending enough. It is more likely that some of the middle value customers have been enticed to spend more. The alternative hypothesis for a 60/20 distribution is the top 20 of companies could be naturally enticed to spend more with the firm. Perhaps they do not because the firm's CVP is undifferentiated against significant rivals. In this case then marketing and manufacturing are still providing differential value because there are some customers who are more prepared to spend. The final hypothesis is the distribution has nothing to do with the actions of the firm but simply reflects the size distribution of customers in the market.

The variation in business distributions does not seem to be limited to marketing activities. Here is a case from *Service Science: Analysis and Improvement of Business Processes* (2020) with a 17/60 distribution¹³⁵:

Consider an apartment broker located in a university town that assists students in finding suitable apartments. The company uses the messaging platform WeChat to interact with its customers. It includes a website with apartment lists and an on-line application. It also arranges visits to apartments for students, develops lease contracts, and arranges for rental payments. The support group of the company responds to questions sent using WeChat. The subject of each message is categorized and tabulated every quarter. A Pareto chart for the least three months is shown in Figure 4.7, which lists the subjects according to their frequency. Its left vertical axis shows the number of messages while the right vertical axis shows the cumulative percentage. The Pareto chart shows that the top 2 of the 12 categories of questions (17% of the categories) accounted for about 60% of the messages. The largest number of WeChat messages concerns the on-line application. This result should motivate the firm to give highest priority to improving this aspect of the service, which could contain confusing entries or terms that many students do not understand. It is notable that the second most popular category is the "other" category. This result is unfortunate because

nothing can be done to reduce these questions without collecting more data on their specifics. It is not uncommon for data collection systems to lack the types of updates necessary to keep the information highly relevant.

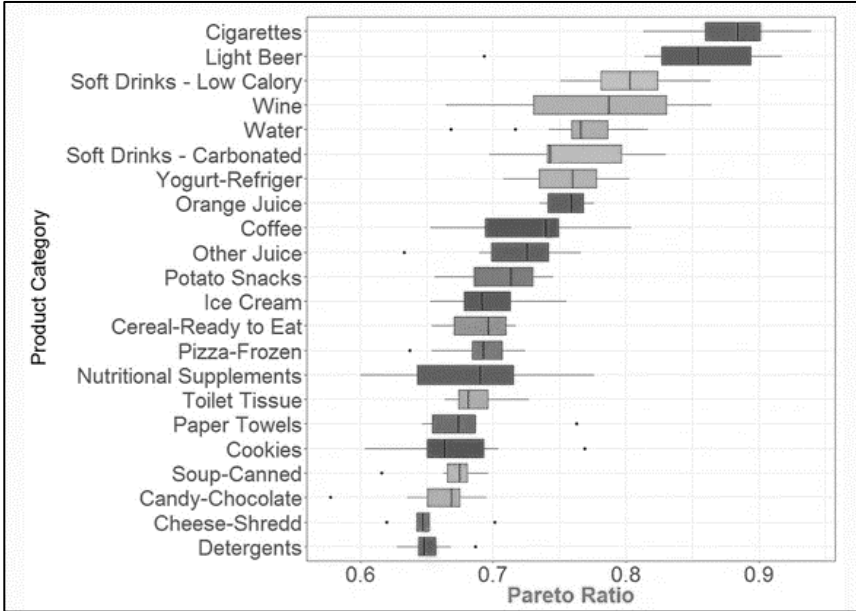
FIGURE 7: PARETO CHART OF WECHAT ACTIVITY



Another good study relates to the Pareto rule for frequently purchased packaged goods using consumer panel dataset from A.C. Nielsen with 6 years of purchase histories from over 100,000 US households. Kim et al found:

...potential factors such as brand attributes, category attributes, and consumer purchase behaviour to explain variation in the Pareto ratio at the brand level across products. Our main conclusion is that the Pareto principle generally holds across a wide variety of CPG categories with the mean Pareto ratio at the brand level across product categories of .73. ...niche brands are more likely to have a higher Pareto ratio. Finally, brand/category size, promotion variables, change-of-pace brands, and market competition variables are negatively correlated with the Pareto ratio¹³⁶.

FIGURE 8: PARETO DISTRIBUTION OF CONSUMER PACKAGED GOODS



Switching from product sales distribution to a study of how consumption of products varies for department store shoppers shows heavy customers have Pareto distributions ranging from 50/20 to 73/20 for consumer packaged goods (CPG)¹³⁷. The overall average CPG Pareto was 58/20. Customer Pareto distributions varied from 38/20 to 67/20 for other product categories.

Pareto applies to other aspects of the customer journey beyond a sales focus. For example Jum'a et al (2023) analysed the time and cost of value added services in warehouse logistics for two processes¹³⁸. The Pareto distributions of cost was 82/20 and 77/20. Rosak-Szyrocka et al (2022) analysed complaints for energy market services and found a 33/20 distribution. These cases will be used to show bottle necks in ToC processes later¹³⁹.

FIGURE 9: PARETO DISTRIBUTION OF DEPARTMENT STORE CUSTOMERS

Study	Year	Analysis Period	Unit of Analysis	Scope	Countries	Average Top 20% Contribution
CPG categories						
Schmittlein, et al	1993	1 year	Transactions	4 CPG categories	US	52%
Jarvis et al	2003	1 year	Transactions	4 wine sub-categories	AU	64%
Sharp and Romaniuk	2007	1 year	Transactions	CPG categories	US, AU, ZA	52%
Romaniuk and Wight	2015	1 year	Transactions	12 CPG categories	UK	61%
Romaniuk and Sharp	2016	1 year	Transactions	9 CPG categories	BR, CN, ID, IN, KE, MX, MY, NG, PH, TR	53%
Graham, et al	2017	1 year	Transactions	22 CPG categories from supermarket loyalty card	UK	50%
Anesbury, et al	2020	1 year	Transactions	5 fresh fruit and vegetable categories	US	61%
Kim, et al	2017	6 years	Sales value	22 CPG categories	US	73%
Graham, et al	2017	5 years	Transactions	22 CPG categories from supermarket loyalty card	UK	60%
Unweighted average (all CPG studies)						58%
Other categories						
Litvin	2000	1 year	Transactions	Vacation travel	SG	40%
Brynjolfsson, et al	2011	1 month	Transactions	Women's clothing – online or catalog	US	57% (<i>catalog</i>) 53% (<i>online</i>)
Bennett	2016	10 years	Transactions	Airplane purchases	Global	56%
McCarthy and Winer	2019	2 years	Transactions	Product and service companies	US	67%
Martin, et al	2020	1 year	Transactions	10 types of retailers	US	50%
Martin, et al	2020	1 year	Value	10 types of retailers	US	38%
Unweighted average (all non-CPG studies)						52%

Variation is Significant

The various cases above show Pareto applies to customer communication content, product sales, customer value, production times, production costs

and complaint causes. This confirms the core concept of Pareto: Some causes are significantly more important than others. Considering averages is less useful than understanding the variations inherent in result distributions. This is because focusing resources to manage for more important causes creates a differential return on investment. The challenge is to work out what are the most important causes. There is a subtlety here worth pointing out because of how AI can be applied to Pareto.

The Pareto distribution is fractal. What this means for an 80/20 distribution is there is another mini 80/20 a level down from the base 20 of most important causes. And there is another level down as well. What this means is if 20% of causes create 80% of results, then 20% of the 20% of the causes creates 80% of the 80% of the results. Also that 20% of the 20% of the 20% of the causes creates 80% of the 80% of the 80% of the results. In simpler words:

- If 20% of causes drive 80% of results [first order cause]
- Then 4% of causes drive 64% of results [second order cause]
- And 0.8% of causes drive 51% of results [third order cause]

The two lower levels are easier to remember as follows. Second order causes are the one in twenty things you to create around two thirds of your results. Third order causes are the 'one percents' create around half your results. Even if the distribution is 60/20 second order causes (one in twenty things) create around a third of your results, and third order causes (one percents) create just over a fifth of your results. Machine learning and AI can be incredibly powerful for identifying second and third order causes from large data sets with many variables.

For example, a study of 39,000 Amazon customers by Alpine.AI and Infoscout showed that Amazon Echo owners increased their spending on Amazon by 29% after purchasing the device¹⁴⁰. In addition, they decreased spending on grocery and general retail products bought elsewhere. The increased spending on Amazon is mainly driven by convenience, with consumer packaged goods being the most purchased items. Voice-mediated retail is growing rapidly, at around 50% per year, and while Amazon is currently the market leader, Google Home is gaining ground by being able to buy from dozens of retailers. This suggests getting people to use Alexa via their Echo devices is a first order cause. Second order causes might include identifying the subset of customers who are likely to use voice interfaces for purchasing and then overlaying the timing, type, and price of specific recommendations to encourage them to purchase. Amazon could focus on getting these customers to take up an Echo if they don't already have one. These customers are more valuable to Amazon if they use Alexa, so it would make sense to gift them one or at least discount the normal purchase price. One percents for voice retail might include analysing each customer's purchase history on

Amazon along with their viewing habits on Prime to determine their personality type. This personality information could then be applied to modify the tone, style, and content of the interactions with each customer to help them experience their version of convenience.

A high-level version of matching style of convenience to personality could be implemented using the DISC framework. The DISC personality framework is a behavioural model that seeks to describe and classify human personality types based on four main categories: Dominant (D), Influential (I), Steady (S), and Conscientious (C). Each type is associated with a particular set of behavioural traits, tendencies, preferences, and communication styles. Individuals who fall into the Dominant category are often described as assertive, decisive, and task focused. They are comfortable taking charge and leading others, and are often attracted to roles that involve control, challenge, and problem-solving. Dominants want control. Individuals with an Influential personality type, on the other hand, are typically outgoing, enthusiastic, and people oriented. They enjoy socializing, networking, and building relationships, and are often drawn to roles that involve creativity, persuasion, and public speaking. Influencers want to have fun. Those with a Steady personality type are typically described as cooperative, patient, and reliable. They prioritize stability, harmony, and teamwork, and are often drawn to roles that involve supporting and helping others. Steady types want things to be easy. Finally, individuals with a Conscientious personality type are typically detail-oriented, analytical, and systematic. They value accuracy, precision, and order, and are often attracted to roles that involve planning, organizing, and problem-solving. Conscientious types want things to be perfect. Understanding these different personality types can be useful in various settings, such as team building, conflict resolution, and leadership development.

Applying this to Alexa suggests different interaction styles for different user personalities because each perceives convenience slightly differently. Alexa might be best to cut to the chase and give three options for Dominant type consumers to be able choose from so they can exercise control. Influencer types might be more interested in either the product with the highest ratings or the most creative and novel. They may also be more interested in 'surprise me' type recommendations. Steady types may find more appeal in offers that automatically set up subscriptions, so they don't have to repeatedly deal with reordering. Conscientious types may prefer to be able to have Alexa select a purchase based on several different variables including prices, volumes sold, ratings, warranty, feature sets and delivery times. They would likely value being able to weight these different variables to improve their purchase decisions.

Companies don't tend to do much of this now with AI because they are still operating on the pre-AI paradigm. This case from Estée Lauder is typical¹⁴¹. The company only used AI to identify a first order cause for consumer engagement with the brand (emphasised by the author in the quote below).

...To increase online brand awareness among young women, the company encouraged potential customers to subscribe to their newsletter and encouraged the audience to leave contact information to receive samples of their products.

The marketing strategy was extended to all devices. In addition, the brand wanted to strongly keep costs of the promotional campaign, requiring the optimization of the cost per click (CPC) and cost per contact (CPL). The Appier, through its CrossX platform, identified the devices owned by individual users and, through the functionality "Lookalike", reached and identified new potential profiles that correspond to the desired profile by Estée Lauder.

*The result was obtained by analysing data stored in the database CrossX, data collected from over 3,000 directly managed campaigns Appier over the years. The use of AI has allowed Estée Lauder to reduce not only the cost of the marketing campaign but also the time of conversion of users into customers. In addition, the brand has been able to discover that **the majority of users**, who have left their information in exchange for free samples offered, **were people who came from ads dedicated to the theme of aesthetics and not so much in skin care.***

...The partnership with Appier allowed to Estée Lauder a 300% increase in the attractiveness of its ads promotional cross-screen, reaching in some cases spiked 1,100%.

...Estée Lauder has obtained the following results:

- *CPC reduced by 43%.*
- *CPL reduced by 63%.*
- *The number of clicks increased by 74%.*
- *The number of conversions increased by 167%.*

The cross-screen conversion devices were considering at least 3:

- *11 times higher than the average for the PG;*
- *4times higher than the average for the tablet;*
- *3times higher than the average for smartphones.*

Just like the Estée Lauder case outlines, promotions on products like Echo tend to be ubiquitously available. User purchase interactions are designed to be the same workflow regardless of the user's personality. The personality segmented example above suggests how to better serve customers based on their personality. It is indicative only. A proper AI would analyse data to identify strategies to maximise consumption on an individual basis. It would also segment consumers into far smaller groups than simply the four in DISC. Ultimately AI with this level of discrimination could treat every single consumer as a unique market segment and work out the one percents that matter to them individually. This is potentially really good and really bad for consumers. Good because they will make better purchases and enjoy a more simple, satisfying, and surprising CX. Bad because controlling their spending will be so much harder!

You can image a scenario where Alexa's AI identifies five days in a row where a customer started watching TV later at night. That might mean longer hours are being spent at work or commuting. The AI could emulate empathy for the customer and determine they really need some help relaxing. After all the customer seems to have earned it after a harder week at work than normal. Cross referencing with past purchase history could improve recommendations on a range of products from aromatherapy and fluffy slippers to a state-of-the-art massage armchair or simply home delivered take away food for dinner. Timing a promotion offer to occur at the exact moment where the consumer's willpower to resist a really good deal is at its lowest maximises the chance of conversion. The AI could use data mining techniques to identify this sweet spot. Finally the offer could be priced at either a discount or premium specific to that consumer to maximise potential margin won.

This may all seem like science fiction. Yet a version of this happened more than a decade ago in 2012. The "Target pregnancy prediction scandal" involved a large retail corporation in the United States¹⁴². Target used data mining techniques to identify customers who were likely to be pregnant based on their purchasing habits. They created a pregnancy prediction model that used customer data such as purchases of unscented lotion, cotton balls, and certain vitamins to predict whether a customer was pregnant. Once Target identified customers who were likely pregnant, they began to send them targeted marketing promotions for baby products, such as diapers and baby clothes.

One customer's teenage daughter began to receive these promotions, which led to a confrontation with the father, who was upset that Target was promoting baby products to his daughter. After investigating the matter the father discovered that his daughter was in fact pregnant. He then realized that Target had correctly predicted his daughter's pregnancy before he was aware of it himself, based on her purchasing habits. The case sparked a debate about the ethics of data mining and the use of personal information for marketing purposes. Target eventually changed its marketing strategy to be less intrusive and to respect customers' privacy. Importantly Target's data mining techniques were powerful enough to predict pregnancy of a customer based only on their purchasing habits with last gen AI. Next gen AI will have access to more data and be able to crunch better conclusions to deploy more effective promotions. It will be able to find the one percents.

Identifying one percents even in 60/20 distributions is an extremely powerful commercial strategy because of the potential to reallocate resources effectively. Taking 10% of a marketing, manufacturing or customer support budget and redirecting it to one percents creates around a 300% growth! Historically human managers have not been able to get much below first order Pareto causes because it is too hard to collect and manage the complexity of data required. In the same way the smartest AI can beat every human on the planet at games like chess and go they will be able to market, manufacture and support us better than humans in many situations.

• Theory of Constraints

Theory of Constraints (TOC) has been adopted by various organizations across different industries to improve their operations and increase efficiency. While many companies don't necessarily advertise their use of TOC, some examples of companies that have been known to apply this management philosophy include:

Procter & Gamble - The consumer goods giant has used TOC to optimize its supply chain, streamline manufacturing processes, and improve overall efficiency.

Intel - The semiconductor manufacturer has applied TOC to address bottlenecks and increase the throughput in its chip manufacturing process.

Tata Steel - One of the world's leading steel producers, Tata Steel has used TOC to improve its operations, minimize inventory, and enhance overall productivity.

Delta Airlines - The airline has employed TOC to optimize its maintenance processes and reduce aircraft downtime, which ultimately leads to better customer service and increased profitability.

Boeing - The aerospace manufacturer has implemented TOC in its production facilities to reduce lead times, increase throughput, and improve overall efficiency.

Amazon - While not explicitly stated, Amazon's focus on customer-centricity, streamlining operations, and continuous improvement aligns well with TOC principles.

Drum-Buffer-Rope

The Drum-Buffer-Rope (DBR) concept is an application of the Theory of Constraints in production planning, aimed at increasing flow and throughput by identifying and leveraging the system constraint¹⁴³. DBR focuses on synchronizing production to the constraint while minimizing inventory and work-in-process. The three essential elements of the DBR concept are:

Drum: The drum represents the constraint or the bottleneck in the system, also known as the weakest link. The speed at which the constraint runs sets the "beat" for the entire process and determines the total throughput.

Buffer: The buffer is used to protect the constraint, ensuring that it always has work to process and does not run out of materials. Buffers are strategically placed before the constraint to absorb variations in the production process, preventing the constraint from being starved and maximizing its efficiency.

Rope: The rope is a communication mechanism that connects the constraint with the rest of the production process. It acts as a signal to "pull" new items of work into the system only when an item is processed and finished moving through the constraint. The rope ensures that the entire production line is synchronized and subordinated to the speed of the constraint, maintaining a smooth and efficient flow of operations.

DBR is used in various industries to improve overall efficiency¹⁴⁴:

- Robert Bowden Inc, a building products supplier, saw a 20% increase in orders processed without adding staff and reduced expediting of orders.
- Harris Corporation, a semiconductor wafer manufacturer, reduced cycle time by 50% and doubled output after implementing DBR.
- A cable and telecommunication equipment manufacturer in Dallas, Texas, reduced work-in-progress to one-third, decreased raw materials inventory value by 30%, and increased on-time completion of jobs by more than 30% after implementing DBR.
- A facility that cuts metallic sheets into smaller coils implemented DBR and reduced lead time from 21-182 days to a stable 10 days, and customer service level increased from 34% to 87%.

- An Italian chemical company producing dyes and pigments implemented DBR, resulting in a decrease in raw materials and finished goods inventory, and an increase in stock turns, which almost doubled between 1999 and 2001.
- Alkco improved its lead time reduced to one week, on-time delivery increased to 98%, sales volume increased by 20%, and before-tax profit increased by 42% over 5 years.
- Trane Co. of Macon, Georgia, doubled its output from three to six units per day with the same workforce after implementing DBR.
- Remploy, a military garment manufacturer, implemented DBR in their two plants, resulting in a 19% increase in Throughput, a 13.4% increase in output per employee, a reduction of more than 50% in WIP, and a 7% decrease in absenteeism, with only a slight increase in transportation costs.

To implement DBR, companies can follow these steps:

1. Identify the constraint: Find the limiting factors within the operations, which can be a physical resource, process, or policy constraint.
2. Exploit the constraint: Make the most efficient use of the constraint by ensuring it is working to its full capacity and is never idle.
3. Subordinate everything to the constraint: Adjust all other processes in the system to work at the same pace as the constraint.
4. Elevate the constraint: Take actions to increase the capacity of the constraint, such as adding resources, training, or investing in new equipment.
5. Repeat the process: Continuously identify and address new constraints as they emerge.

By effectively implementing DBR, companies can optimize their production processes and minimize bottlenecks. As a result, they may discover new ways to meet customer needs more effectively. For example, improved production efficiency may enable a company to offer faster delivery times, better quality products, or lower prices compared to competitors, all of which can be attractive to customers. These can form the basis of a compelling 'Mafia Offer'.

The Offer You Can't Refuse

A Mafia Offer in the context of the Theory of Constraints (TOC) is an unbelievably good offer that a company makes to its customers. The attractiveness of the offer makes it impossible for the customers to decline and difficult for the competitors to match. Mafia Offers are not a positioning or a tag line. They can only be created by satisfying a significant need of the

market to the extent that no other significant competitor can. Here are some examples of Mafia Offers:

Hyundai's Mafia Offer was to buy back customers' new cars if they were laid off within two years of purchase to increase sales. Domino's Pizza famously promised to deliver pizzas in 30 minutes or less, or the customer would get their pizza for free to differentiate against rival pizza chains. Xerox offered copies when and where customers needed them, at a fixed rate per page (instead of selling or leasing copiers) to compete against non-consumption. Liberty Bell Telecom implemented a mafia offer and saw significant improvements in their sales process, tripling new customer sales and increasing monthly recurring revenue growth from \$25,000 to \$70,000¹⁴⁵. The basis of their offer was using DBR to lift their customer service. It was so good Liberty Bell experienced less than one percent customer churn and bad debt per month after executing¹⁴⁶.

None of the above customers have built a Mafia Offer on AI. Cogito is a company that uses emotion and conversation AI to deliver real-time coaching and guidance to contact centres and frontline teams¹⁴⁷. By detecting and interpreting human emotional signals through text and voice, Cogito's Emotion AI can provide real-time coaching for contact centre agents, allowing them to show empathy and improve their customer interactions.

The AI coach works by analysing hundreds of behaviours within the conversation and providing insights and guidance in real-time to improve the interaction between the agent and customer. Cogito's AI Coaching System provides personalized coaching that drives increased performance, making top performers more consistent and improving the behaviour of weaker performers. This benefits both the agent and the customer, leading to more productive and empathic interactions over time. This leads to improved first call resolution metrics and customer satisfaction scores in call centres where Cogito has been implemented.

An article from Time discusses the use of Cogito¹⁴⁸. Cogito was rolled out to 10 call centres at MetLife in 2019. Managers said it improved first call resolution metrics by 3.5% and customer satisfaction by 13%. Additionally, the program helps agents (who take an average of 700 calls a week) have more "human" conversations. One employee said Cogito helped her cut her average call time nearly in half, while another said that it helped her slow down when she was speaking.

The article also discusses the potential for AI to displace as many as hundreds of millions of workers over the next decade or two. Cogito's scientists say that a system that replaces human workers is still at least a decade away. Lastly, the article touches on the potential long-term effects of spending eight hours a day, five days a week under AI direction, suggesting that it may have effects

beyond how someone speaks at work. One Cogito user notes that after using the program for about a year, her boyfriend began noticing a difference in the way she spoke at home, indicating that her communication style had become more direct in all conversations.

Normally Cogito's kind of real time coaching can only be delivered as a one-to-one interaction between a human coach and each frontline team member. Cogito's Mafia Offer is to be individually available for each team member for their entire work week at a fraction of the cost of a single coach.

Mafia Offers in Healthcare CX

The above Mafia Offers were described by value they created for the companies bringing them to market. The flip side in CX is to be customer centric and consider the additional value unlocked for customers. Customer centric firms share in this value. It is the reward for innovating the customer experience. Two healthcare examples are provided to illustrate this concept. There are differences with healthcare contexts important to call out first.

Healthcare is normally better thought of as a system than a market. In a market suppliers and customers tend to have free choices about what to buy and sell for what price. Suppliers can use marketing, new product development and improved CX to earn more for their offers. Consumers can compare different suppliers to get the best deal for themselves. Markets work best where supplier competition is effective and customers can make informed choices. Most markets operate within some reasonable regulatory limits to try and ensure market failure is avoided. Both suppliers and consumers are more constrained in healthcare than in other markets.

Some of the constraints for suppliers include regulations, government funding and the relative scarcity of properly trained healthcare professions. The most important constraint for consumers is it is much harder to properly advocate for yourself if you are sick, injured or dying and need care. Even if it might be rational and economically efficient to ask someone who is dying to pay everything they have for a cure, few people are ok with this.

Healthcare is a crucial aspect of society, and as such, it is supported by taxes and insurance premiums. The funding for healthcare in Australia comes from both the Australian and state/territory governments, as well as the non-government sector. The Australian government funds most of the spending for medical services and subsidized medicines, while state and territory governments fund most of the spending for community health services. While healthcare is important for all Australians, the burden of funding it is disproportionately carried by the rich and the young. In 2016-17, \$1 in every \$10 spent in Australia went to health, and the government usually funds most of the spending for medical services and subsidized medicines¹⁴⁹. This means

that the taxes paid by high-income earners are a significant source of funding for the healthcare system. Additionally, the elderly tend to require more healthcare services and support than younger individuals. As such, they may receive more benefits from the healthcare system, which can be subsidized by the taxes paid by younger individuals. It is worth noting that this is not a simple or straightforward relationship, as many factors contribute to healthcare costs, including lifestyle, genetics, and the environment.

Inequalities can exist within healthcare systems. For example, a 2017 US government report highlighted that making healthcare services more available does not mean that they will be equally accessible to all ethnic and socioeconomic groups¹⁵⁰. Additionally, some argue that the current US healthcare system tends to delay or deny high-quality care to those who are most in need of it but can least afford it, contributing to healthcare disparities for disadvantaged groups¹⁵¹.

Health insurance has become a common means for individuals to manage the high cost of healthcare. From an economic efficiency perspective, insurance provides several benefits and disadvantages. One of the advantages of health insurance is that it can increase access to healthcare services, which can improve health outcomes and productivity in the workforce. Insurance also helps to spread the financial risk of healthcare costs across a larger pool of individuals, which can make healthcare more affordable for everyone. Moreover, it incentivizes people to seek preventive care and early treatment, which can help reduce healthcare costs in the long run.

However, health insurance can also have several drawbacks. One of the main disadvantages is that it can distort the cost of healthcare, leading to overutilization of healthcare services. Providers often do not have to ask customers for their full schedule because this is mediated by insurance. This shifts the consumer away from total cost to merely considering the out-of-pocket costs when choosing to access health care. As a result, consumers may choose to access more healthcare services than they would if they had to pay the full cost out-of-pocket, which can drive up healthcare costs and lead to waste and inefficiency in the healthcare system. Another disadvantage of health insurance is that it can create moral hazard. Because insurance covers a large portion of healthcare costs, consumers may not have the same incentive to shop around for the best value or to demand the most cost-effective treatments. Providers may also have less incentive to keep costs low or to provide high-quality care, as they are insulated from the true cost of care by insurance. Furthermore, health insurance can be a complex product that may not always be well understood by consumers, leading to confusion and mistrust. Insurance can also be a source of administrative burden and costs, as insurers and providers must spend resources on billing, claims processing, and other administration.

The analysis above does not change that CX is important in healthcare because it is a type of market. Providers and insurers are still often in competitive situations and consumers do exercise choice on both a rational and emotional basis. This is true even if both sides are somewhat more constrained than similar agents in other verticals.

Grand Rounds is a healthcare service company that aims to provide high-quality medical care by connecting patients with top medical specialists for second opinions and expert medical advice¹⁵². They work with employers and health plans to cover the cost of their service, which often results in patients receiving second opinions at no out-of-pocket expense. Their service focuses on offering personalized and convenient access to high-quality healthcare. They maintain a network of highly qualified physicians and specialists, carefully selected based on their expertise, experience, and performance. This network allows them to match patients with the most suitable specialist for their specific condition or medical issue.

To access Grand Rounds' services, patients must be part of an organization that partners with Grand Rounds, such as an employer or a health plan. If a patient's organization is partnered with Grand Rounds, they can register on the Grand Rounds website and request a second opinion. The patient will be asked to provide relevant medical records, and a care coordinator will help guide them through the process.

Once the patient's information is submitted, Grand Rounds will match the patient with an appropriate specialist from their network. The specialist will review the case and provide a detailed, written second opinion, including recommendations for treatment options. This second opinion can be shared with the patient's primary care physician or specialist, who can use it to inform their treatment decisions.

In addition to second opinions, Grand Rounds also offers other services like care coordination, assistance in finding in-network doctors, and support for navigating the healthcare system. Their goal is to help patients make informed decisions and receive the best possible care. Grand Rounds second opinion service is an example of CX innovation in healthcare without utilizing either last generation or next generation AI. A free second opinion is a potential Mafia Offer for the firm.

In contrast Babylon Health is a digital-first healthcare provider that uses an AI-powered platform to connect patients with healthcare professionals¹⁵³. They aim to revolutionize healthcare delivery by providing preventative healthcare instead of reactive sick care through the devices people already own. Babylon Health's business model is based on a subscription-based private healthcare service that was launched in the UK in 2013, which has now expanded globally, serving over 24 million people across the world.

Babylon Health's virtual care service, Babylon 360, combines AI with human medical experts and emphasizes virtual care through apps, chats, and video consultations, aiming to keep patients out of the hospital and keep costs low. Additionally, Babylon Health is active in sixteen countries on four continents and offers both an AI-based symptom triage system that works for worried millennials and a risk-bearing care model. The company's business model seems to be more of a Mafia Offer than Grand Rounds. Here are some key aspects of Babylon Health's business model supporting this assertion:

Accessibility: Babylon Health makes it easy for patients to access healthcare professionals from anywhere, anytime, through their smartphones or computers almost on demand. This convenience is especially appealing for those who have limited access to healthcare providers or who face long waiting times for appointments. In CX terms it provides increased utility at the Use and Support stages of the patient (customer) journey. It may also contribute to the Awareness, Consideration, Repeat and Referral stages.

Affordability: Babylon Health offers its services at a lower cost than traditional healthcare options, making it more affordable for a broader range of people. Some services are even provided free of charge, depending on the user's location and healthcare plan. This is possible because the AI allows face to face time between patients and doctors to be more efficient. Affordability is a critical aspect for Consideration, Decision, Purchase and Use stages of the patient (customer) journey. Again it may impact Awareness, Repeat and Referral stages as well.

AI-driven services: Babylon Health uses artificial intelligence to power its symptom checker and triage services. The AI can provide medical information and recommendations based on the user's symptoms, making it easier for patients to understand their health issues and identify the most appropriate course of action. Automating part of the health care experience without removing the interaction with health care professions and changing the paradigm to be preventative as well as reactive potentially lifts CX for all patient (customer) journey stages. This element has most relevance for the Use, Support and Repeat stages because AI makes the experience of going to the doctor both simpler and potentially offers a more informed diagnosis. Better diagnoses are more satisfying and potentially surprising.

Integration with healthcare systems: Babylon Health partners with national health services and insurance providers, ensuring that their services are integrated with existing healthcare systems. This allows patients to receive seamless care and access additional services if needed. Most relevant to the Purchase and Use stage of the customer journey.

There are aspects of the Babylon Health offer that are traded off against the various CX improvements listed above. Patients (customers) are not certain

they will always see their same doctor. They do not have an option to avoid the AI triage to only speak with a GP about their health. This is a significant negative for those who need more emotional support as well as expert, rational healthcare from their GP. Traditionally the doctor patient relationship has been a foundation aspect of primary care in many countries. Babylon Health's trade-offs tend to have the potential to polarise consumers towards the firm and its offering. Polarisation is exactly what good segmentation should do. This makes the company's business model disruptive to traditional healthcare systems. It is Mafia Offer because it does provide a compelling offer to some patients (customers) that is hard for other healthcare organisations to copy because they are not founded on an AI triage system or a subscription model.

• Building the Road Map


To build a CX roadmap using Pareto Analysis and the Theory of Constraints with AI follow these steps:

1. **Document the Current State of CX:** Establish a baseline understanding of existing customer experiences and the internal efforts underway to track and improve those experiences.
2. **Conduct Pareto Analysis:** Identify the most important causes of good results across the customer journey by analysing data and identifying the 20% of factors that contribute to 80% of the positive outcomes. These outcomes can be any of the 5 Ps (portion, premium, permanence, promotion or pull)
3. **Apply AI-driven Nested Pareto Analysis:** Use AI to further refine the Pareto Analysis by identifying the 20% of factors within the initial 20% that contribute to 64% of the positive outcomes (64/4). Then, apply the analysis again to this 4% to find the factors that account for 51% of the positive outcomes (51/0.8). This reveals the one percents that really matter.
4. **Create a Customer Journey Map:** Outline all the opportunities, pain points, and interactions across the customer journey, which can be used to guide your CX improvements. There should be both a starting journey and a future version detailing the target state experience because this makes it possible to communicate to internal stakeholders what is required and how each of them relates to the planned changes.
5. **Identify Constraints:** Using the Theory of Constraints, pinpoint the bottlenecks or limitations in your current processes that hinder the achievement of the desired customer experience. Constraints can be due to limited capacity, poor quality, high costs or low customer centricity.

6. **Apply Drum Buffer Rope:** Based on the identified constraints, redesign your business workflows by synchronizing processes, allocating resources, and setting priorities. The Drum represents the constraint, the Buffer protects the constraint from variability, and the Rope ensures that the constraint is always working at maximum capacity. Use AI to automate key bottlenecks and release constraints without reducing customer empathy.
7. **Define Your CX Vision:** Establish a clear vision of the desired customer experience, which should align with your overall business goals and objectives. This should also link to the firm's financial and competitive objectives. The vision should be costed to forecast an expected return on investment. AI can be used to improve revenue gain and cost reduction modelling.
8. **Set CX Metrics and Goals:** Define measurable objectives and key performance indicators (KPIs) to track the progress of your CX roadmap implementation. Data collection should be automated where possible.
9. **Implement Changes:** Execute the process redesigns and improvements based on the findings from the nested Pareto Analysis and DBR process redesign. Additional Pareto causes may be discovered, and ToC bottlenecks should be updated as changes are executed.
10. **Monitor and Optimize:** Continuously monitor the performance of your CX strategy, gather customer feedback, and make data-driven adjustments to improve customer experiences. Ensure your voice of customer program can track CES, SAT or NPS at key touchpoints.
11. **Review and Iterate:** Regularly review your CX roadmap, making necessary adjustments to ensure alignment with changing business goals, customer expectations, and market conditions. Expect your CX uplift will cause a reaction from market rivals at some point.

The above process represents a source of new competitive advantages for firms who are first movers to utilise next gen AI for CX uplift. Eventually something like the above process may become more or less best practice. Similar to how Six Sigma remains a highly regarded set of quant-based problem-solving tools currently and quality management was an important philosophy in the 1990s. The next section expands on the ideas in this chapter to show how to make these approaches hard to copy.

VII. The New Sources of Competitive Advantage



AI was determined to win at all costs.

Business strategy involves the identification, creation, and exploitation of competitive advantage to achieve commercial objectives. Mintzberg's 5Ps model of strategy represent different ways to understand strategy. The 5Ps are Plan, Ploy, Pattern, Position and Perspective¹⁵⁴.

Plan: Strategy as a plan involves developing a deliberate course of action in advance with a clear purpose and objectives. It focuses on setting goals, formulating initiatives, and outlining the steps necessary to achieve them. The plan provides a roadmap for the organization's future direction. Exploiting AI capabilities can be part of a firm's CX strategy to lower cost or increase revenue as previous chapters have outlined. Plans are powerful for creating alignment and ensuring a strategic approach has been thought through. However strategic plans are limited because no plan survives contact with reality during execution.

Ploy: Strategy as a ploy refers to using tactics and manoeuvres to outsmart or gain an advantage over competitors. It involves making calculated moves and deploying tactics to influence the competitive landscape, disrupt rivals, or create favourable market conditions. Ploys are generally designed to take advantage of rival sunk costs or policy decisions. They are of limited effectiveness against market disruptions, indirect and new entrant rivals. ConXtech, a modular-construction company based in California, utilizes AI to gain control over the bidding process in construction projects. During the project-development phase when ConXtech is solicited by owners and developers they employ AI to assess the viability of the project and make informed decisions regarding bidding. They may initially suggest a reticence to bid on a project until their AI suggests the situation is optimal. In this context, AI may be eroding the experience of the owners and developers who want to leverage competitive bids to get better pricing.

Pattern: Strategy as a pattern emphasizes the importance of understanding and leveraging patterns or trends that emerge from past actions and decisions. It involves recognizing recurrent behaviours or actions that have proven successful and using them as a guide for future strategies. Patterns are effective to the extent they are generalisable. Taleb's turkey problem from *The Black Swan* (2011) serves as a metaphor for the limitations of relying on past observations to predict the future¹⁵⁵. The turkey is fed every day, leading it to believe that this pattern will continue indefinitely, until the unexpected event of Thanksgiving occurs. The key lesson is that relying solely on historical data or patterns to forecast the future can be misleading and dangerous. The current state of AI and way software automation has traditionally been utilized is the subject of this entire book. LLMs are analogous to the unexpected Thanksgiving event for the turkey. What is interesting is it is not clear who is getting to hold the axe in this scenario. In many markets it is still up for grabs.

Position: Strategy as a position focuses on the organization's relative position in the market or industry. It involves identifying a unique and favourable market position that sets the organization apart from competitors. Position-based strategy may include aspects such as differentiation, cost leadership, or niche targeting. Positioning involves an inherent paradox: ideally a firm wants to occupy an attractive market position with high demand volumes, but this will likely be crowded with rivals. The alternative is to find a market position where it is the only supplier, but this will likely be subject to low demand volumes. Positioning is the foundation of Kim and Mauborgne's *Blue Ocean Strategy* (2015)¹⁵⁶. They refer to red oceans as existing markets that are highly competitive and saturated, where companies strive to outperform rivals for a share of the existing demand. Blue oceans represent unexplored market spaces with untapped potential, where companies create new demand by simultaneously pursuing differentiation and low cost, making competition irrelevant. In CX, AI offers the possibility to both lower costs and increase differentiation simultaneously.

Perspective: Strategy as a perspective emphasizes the need to develop a collective mindset or worldview within the organization. It involves considering different viewpoints, values, and beliefs to shape the strategic direction. Perspective-based strategy recognizes that diverse perspectives can lead to more comprehensive and innovative strategies. Customer centricity can be understood as an example of the Perspective strategy from Mintzberg's 5 Ps. It involves aligning the entire organization towards understanding and fulfilling the needs of customers. By adopting a customer-centric perspective, companies prioritize customer satisfaction, tailor their products or services to customer preferences, and strive to build strong, long-lasting relationships with customers. The challenge with perspective-based strategies is the difficulty in coming to an agreement about the specific perspective offering competitive advantage. The basic perspective relevant to AI is whether it is an opportunity, threat, both or neither because all four viewpoints are potentially valid.

The different possibilities about what competitive strategy is and the inherent problems with each approach help explain why effective strategy can be hard to formulate. Two approaches are worth considering the Resource Based View of the firm (RBV) and Dynamic Capabilities View (DCV)¹⁵⁷. These approaches can be considered as Patterns or Perspectives in Mintzberg's 5Ps.

- **Resource-Based View of Strategy**

The first step in formulating business strategy with the RBV of the firm is the identification of potential competitive advantage. This involves understanding the relative strengths and weaknesses of the organisation in its market context. It is a view from inside out. By evaluating the resources a firm has,

organizations can determine their unique position and the potential areas where they can outperform competitors. Competitive advantage can be derived from various sources. Sometimes tangible assets offer an edge. Other times intangible assets underpin the market advantage. Importantly the RBV focuses on the difference between the firm and its rivals.

Many firms may not initially possess a competitive advantage, but it is possible to develop or acquire them through strategic investments. They can be made or bought. Competitive advantage can be achieved by making targeted investments in areas such as technology, research and development, talent acquisition and development, operational efficiency, customer experience, branding, and market positioning. These investments should align with the firm's overall strategic objectives and consider market dynamics, customer needs, and industry trends.

Finally, business strategy focuses on the exploitation of the identified competitive advantage to achieve the desired outcomes. This involves implementing the formulated strategies, monitoring progress, and making necessary adjustments along the way. It requires effective execution, coordination, and alignment of various organizational functions and activities to ensure the strategies are successfully implemented. By continuously evaluating the competitive landscape, adapting to changes, and capitalizing on the unique strengths of the organization, sustainable competitive advantage can be maintained and exploited for long-term success.

The RBV emphasizes that a firm's performance and competitive advantage are determined by the resources it controls. It asserts that certain resources can lead to sustained competitive advantages, which can be summarized using the acronym PROFIT, representing six different resource categories.

Physical Resources: These include tangible assets such as property, manufacturing facilities, equipment, and distribution networks. Physical resources can offer a competitive advantage by enhancing operational efficiency, improving product quality, or reducing costs. The value of mine site is related to size, purity, and ease of access to the ore body it extracts. Physical resources can also offer a competitive advantage where they increase revenues due to proximity to customers or because they draw customers. Hotels in areas of natural beauty with epic views enjoy higher occupancy and revenues than the same hotels located in an airport precinct with no natural beauty. AI can play a part in forecasting the likely future value of physical resources.

Reputation Resources: Reputation refers to the positive perception and trust that stakeholders have in a firm. It encompasses factors like brand reputation, customer loyalty, and corporate image. A strong reputation can provide a competitive advantage by attracting customers, retaining talent, and

enhancing the firm's credibility. This is often bound up in a firm's brand. It is hard to conceive of examples of brands that create lower costs for their firms, but most luxury B2C brands would claim a revenue, margin and volume advantage from their name, logo, and reputation. IBM was famously able to leverage a brand based on reliability in the days where "nobody ever got fired from buying IBM"¹⁵⁸. Firms who intend to bundle AI capabilities within their customer value proposition can benefit from a reputation for reliability, technical superiority or inside knowledge. However for the most part AI seems to be levelling the playing field for many firms because it can lower transaction costs so significantly.

Organizational Resources: Organizational resources include a firm's structure, culture, and capabilities. These resources are related to the firm's internal processes, managerial expertise, and coordination mechanisms. Effective organizational resources can lead to improved innovation, efficient decision-making, and flexibility, thereby contributing to a competitive advantage. The ability to leverage AI is a key organizational resource.

Financial Resources: Financial resources involve a firm's capital, funding, and financial capabilities. Adequate financial resources enable a firm to invest in research and development, expand operations, acquire other companies, or pursue new market opportunities. Having access to financial resources can provide a competitive advantage by facilitating growth and strategic investments. Generally, the nature of financial markets tends to compete any advantages of this category away because both internal and external financiers peg finance costs and availability to systemic risk. AI may offer advantages here for more accurate or lower cost risk evaluations.

Individual Resources: Individual resources refer to the knowledge, skills, and expertise of employees within the firm. These resources encompass the human capital, intellectual property, and experience possessed by individuals within the organization. Leveraging individual resources can lead to innovation, superior problem-solving, and the ability to adapt to changing market conditions. There are many examples of individuals who seem to be crucial to their firm's success. Jeff Bezos founded Amazon and proved to have a remarkable capability to attract investors and not have to pay dividends for many years. Geoffrey Dyson was able to build a company based on re-engineering existing products involved with moving air. Ricardo Semler grew Semco by implementing HR approaches that were counter to best practice. In the field of AI Greg Brockman, Ilya Sutskever, Wojciech Zaremba, and John Schulman are all highly accomplished researchers who founded OpenAI (of ChatGPT fame).

Technology Resources: Technology resources encompass the technological infrastructure, patents, proprietary software, and know-how possessed by a

firm. Advancements in technology can significantly impact a firm's competitive position, productivity, and efficiency. Being at the forefront of technology can provide a competitive advantage through product innovation, process optimization, and operational excellence. Right now the ability to leverage AI technology offers a potential competitive advantage. The question is how valuable AI is.

Merely having resources in different categories does not confer competitive advantage. And analysing for competitive advantage does not simply result in a yes or a no. Different level of competitive advantage can be explained in terms of the degree to which they provide an EDGE. This stand for Extra, Differentiated, Growth and Enduring: Resources that offer an Extra competitive advantage contribute to increased revenues or decreased costs, thus enhancing the firm's profitability and performance.

Resources that are rare and unique compared to competitors provide an initial competitive advantage. Being Differentiated means there are few rivals who can mimic the firm's strategy by leveraging the same class of resources. Historically firms have pursued monopoly positions. More recently oligopoly positions have come into more favour. The profitability of pursuing a monopoly or oligopoly market position is influenced by external regulators who monitor and regulate these market structures. In a monopoly market, the firm can set prices and quantity levels to maximize profits, but regulatory constraints can limit their ability to do so. In an oligopoly market, firms face competition from each other, and external regulators aim to promote competition and prevent anti-competitive practices. Generally regulatory oversight is less in oligopoly markets. In either case a firm can be assumed to have a Differentiated resource.

Some resources are limited to niche opportunities unless they can be scaled effectively. Resources with scalability potential enable a firm to pursue growth opportunities, expand into new markets, and capture a larger market share. The easiest way to conceptualize a growth resource is whether or not its value can compound. One scenario is if the firm can apply the resource for advantage beyond its current niche like Siemens did when it applied engineering capabilities to new products and markets. Another scenario is where the resource gets more valuable as the firm scales. This could happen via increasing brand awareness for reputation resources or through network effects for organizational resources. The final scenario is if the firm can reinvest its above average profits to increase the impact of a resource. When firms are much larger than their rivals in some online market categories, they have the option of bidding up advertising prices to force their rivals out of the market for lead generation. Another is example is where property owners generate equity in their land holdings and use these to secure additional

funds for expanding their portfolio to create local monopolies for lease space in retail precincts.

The final important assessment for a resource is whether it is Enduring. Advantages from resources are sustainable when they are difficult to copy or substitute by competitors. These resources can create barriers to entry and long-term sustainable success. There are three ways a resource can be Enduring. The most common is to exploit a first mover advantage. This is where a firm changes the market for every other firm that goes after them. Patents are designed to do this. By going first and patenting a product or process a firm enjoys the monopoly right to sell for up to twenty years. Similarly, Walmart changed the value of regional retail locations in the US because it bought sites based on their value as farmland before any other retailers. The sellers of the land did not know how valuable the sites were for this different usage, so they sold the land to Walmart based on their potential crop yields. Later rivals to Walmart had to pay more because the owners had figured out the value of their land as retail space because of Walmart's success. The second way to get an Enduring resource is to resolve complexity. Firms who use ToC or Six Sigma to improve their process are essentially resolving complexity. This is inherently hard to copy because the details of the solution are not generally obvious outside the firm. The final basis for Enduring is causal ambiguity.

The Beatles faced multiple rejections from record companies before finally securing their first album contract. It is estimated that they were rejected approximately 20 times. One significant rejection came from Decca Records, who famously turned them down after an audition in 1962, stating that "guitar groups are on the way out". Despite these initial setbacks, the Beatles eventually found success when they were signed by Parlophone, a subsidiary of EMI, and went on to become one of the most successful and influential bands in music history¹⁵⁹. It's not exactly clear what made the Beatles first album such a hit. Their success was obvious only in hindsight. It was much easier to get recording contracts after their debut album when the potential appeal of their music was easier to value. After the breakup of the Beatles, each member went on to have lesser degrees of success in their solo careers. The fact they never reached the same level of critical and commercial success as they had as a group suggests what caused their success was never really made explicit. The cause was at least partly ambiguous.

In summary, the Resource-Based View suggests that firms can achieve sustained competitive advantages by leveraging the six resource categories: Physical, Reputation, Organisation, Financial, Individual, and Technology. These resources offer a potential competitive advantage if they increase revenues or decrease costs (Extra), are rare and differentiated (Differentiated), can be scaled for growth (Growth), and are difficult to copy (Enduring).

Open Source Issues

Most AI deployments to date have not offered a sustainable competitive advantage because they have been easy to copy. This is primarily due to the widespread availability of open source AI frameworks and tools that have democratized access to AI technology. As a result, many companies can readily adopt and implement similar AI capabilities, eroding the uniqueness and exclusivity of any single AI deployment. Here are the reasons behind this lack of sustainable advantage and the implications of the open source approach.

Availability of Open Source AI Frameworks: Open source AI frameworks, such as TensorFlow and PyTorch, have played a significant role in accelerating AI adoption. These frameworks provide a foundation for developing AI models and algorithms, allowing developers worldwide to collaborate, contribute, and leverage a vast pool of shared knowledge and code. The availability of these frameworks has levelled the playing field, making AI technology accessible to a wide range of organizations.

Replicability and Reproducibility: Open source AI frameworks enable the replication and reproduction of AI models and algorithms. Organizations can access and implement similar AI solutions, often without significant barriers. While training large-scale AI models may require substantial computational resources, the fundamental building blocks and techniques are openly available, facilitating the reproduction of similar AI capabilities.

Lack of Proprietary Differentiation: The ease of access to open source AI frameworks means that the underlying technology is not a proprietary advantage for most organizations. This leads to a situation where multiple companies can develop similar AI models or applications using the same foundational tools. Consequently, AI deployments become commodities rather than sources of sustainable competitive advantage.

Rapid Technological Advancement: The field of AI is advancing rapidly, with new techniques and models constantly emerging. The open source community actively contributes to this progress by sharing research papers, code implementations, and novel methodologies. As a result, innovations that initially provide a competitive edge quickly become commonplace as they are assimilated into the open source ecosystem.

The open source approach adopted by many AI developers suggests a different approach is required to leverage AI as a resource for sustainable competitive advantage.

• Dynamic Capability View

Nelson and Winter provide a view extending beyond the RBV in *An Evolutionary Theory of Economic Change (1982)*¹⁶⁰. They highlight the importance of dealing with change in strategic management. The book argues for an evolutionary perspective on economic change, drawing inspiration from biological evolution. Nelson and Winter assert firms and industries evolve over time, responding to changing market conditions and adapting their strategies and capabilities. This underscores the importance of the Dynamic Capability View (DCV). The strategic focus of the DCV is competitive survival in response to rapidly changing contemporary business conditions. It refers to a firm's ability to integrate, build, and reconfigure its resources and competencies to respond to market shifts and create competitive advantages. In one sense they are suggesting the only competitive advantage of lasting value is the ability to innovate (an organisational resource in the PROFIT model). Innovation drives heterogeneity and variation in both firm and industry evolutions. Firms within an industry differ in terms of their capabilities, strategies, and routines, leading to variations in performance and outcomes. This heterogeneity creates opportunities for some firms to adapt successfully and outperform others.

Path dependence is central to the Nelson and Winter's version of the DCV. Path dependence suggests that historical events and early choices can shape future trajectories and outcomes. The choices made by firms in the past can lock them into specific paths, influencing their subsequent behaviour and limiting their ability to adapt to new circumstances. The most obvious example of this is when sunk costs prevent a firm from adopting new technologies or exploiting emerging market opportunities. Feedback mechanisms amplify the effects of early strategic choices in shaping economic change. These loops influence the accumulation and utilization of knowledge, innovation, and learning inside firms and markets. These feedback mechanisms can either reinforce existing patterns or enable shifts in economic behaviour.

For example the German multinational conglomerate Siemens began as a manufacturer of telegraph installations and electrical equipment. Over time they diversified their operations to include various sectors such as energy, healthcare, transportation, and industrial automation. Siemens has consistently been involved in technological advancements, including the development of advanced medical devices, automation systems, and industrial solutions. They have made strategic acquisitions to enhance their technological capabilities. This diversification allowed Siemens to leverage their core capabilities while exploring new markets and industries.

Siemens has a long history marked by research and development. One of Siemens' early innovations was the invention of a telegraph system that used a needle to point to letters instead of Morse code. This approach set them apart in the telegraph industry. Siemens ventured into electric technologies, becoming one of the first companies to build electric locomotives. They also played a significant role in the development of electric street lighting and participated in the early experimentation with AC networks. The common strategic approach in these cases was fast followership in an established market. Differentiation was achieved through technological innovation. Essentially Siemens has tried to leapfrog incumbent competitors in different markets. Their differentiation has been based on engineering capabilities rather than marketing, finance, logistics or human resources capabilities. Siemens continued focus on technological differentiation suggests path dependence. Their approach has been broadly successful over the long term.

Siemens established branch offices and expanded its operations globally. They opened their first office in London in the 1850s and later expanded to Japan and other countries, becoming an international company. Siemens underwent various mergers and acquisitions to expand its business and diversify its portfolio. These included mergers with Schuckert & Co., Nixdorf Computer AG, and various other companies, as well as the acquisition of companies like Allis-Chalmers and Westinghouse Power Generation. This suggests the firm developed an additional competency in acquiring and integrating other firms. This is an example of dynamic capability development because this financial competency goes beyond their core engineering innovation capabilities.

Market signals have also been important to Siemens strategic choices. Siemens has shown a commitment to sustainability by investing in green technologies and solutions. The firm recognized the growing importance of renewable energy and made significant investments in wind power. They built offshore wind turbine manufacturing facilities in the UK and acquired companies in the renewable energy sector. They have focused on developing energy-efficient products and implementing environmentally friendly practices in their operations. Negative feedback has also played a part in the company's evolution.

Siemens faced a major bribery scandal beginning in 2005, which brought significant negative attention and legal consequences. The company was found to have engaged in widespread bribery practices, leading to substantial fines and a complete overhaul of their compliance policies.

What Siemens has not done is exploit Web 2.0. Web 2.0 strategies involve businesses creating online spaces where people come together to share ideas and experiences¹⁶¹. These interactions can result in value creation that

surpasses the sum of the individual parts, with customers themselves contributing to the growth and development of the business.

Flickr is an online photo management and sharing application that allows users to display their favourite photos and videos to the world¹⁶². It also allows users to show content securely and privately to their friends and family. The platform gave users the opportunity to organize photos in their own style, share their photos with friends and family privately or publicly, and create a group for special occasions such as reunions, weddings, and others. Despite facing challenges in its history (including changes in ownership and financial difficulties) Flickr's user-centric features and community-focused approach have allowed it to remain a significant platform for photographers and photo enthusiasts to create and share value.

One of the key features of Flickr is the ability to create 'Sets'. These are collections of images determined by the user. This allows users to categorize their images based on their personal preferences. It makes it easier for them to manage their content and for others to browse it. Flickr also introduced the concept of a 'Photostream'. This is a steady stream of every photo a user has ever shared with the newest at the top. This gives users an easy and convenient way to display their photographic journey to others. It adds a sense of chronology and narrative to their collections. The value created by Flickr for its users goes beyond simply being a space to store and share photos. It's a platform that encourages creativity and community. Flickr allows users to interact with each other, form groups around common interests, and even contribute to a larger pool of publicly accessible photography. This ability to connect with and contribute to a community of photographers and enthusiasts significantly enhances the value users derive from the service.

Firm Evolution

The cases of Siemens and Flickr show they are two very different companies operating in distinct sectors with different purposes and offerings. They have evolved to be different. Siemens AG is the largest engineering company in Europe. Siemens is involved in areas such as industrial automation, energy generation and distribution, medical technology, and transportation. Siemens operates in a B2B (business-to-business) model, providing solutions and services to other businesses and organizations. On the other hand, Flickr provides a platform for users to upload, share, and organize their photos and videos. The platform also offers some social features allowing users to interact with each other's content. Flickr operates on a freemium business model. This means it offers basic features for free with additional features available through paid plans. Flickr operates primarily in a B2C model by providing services directly to individual users.

While both companies provide value to their users, the nature of the value is different due to the differing nature of their operations. Siemens provides technical and infrastructure solutions, while Flickr provides a digital platform for photo storage and sharing. Siemens' value is often measured in improved efficiency, cost savings, or technical capabilities for businesses, while Flickr's value is in the ability to store, organize, share, and interact with digital photos and videos. Importantly Siemens seems to have mastered the DCV approach to competitive strategy. Flickr does not seem to have achieved the same level of dynamism and this has potentially contributed to its financial challenges.

Niche Evolution

In contrast Snapchat, a social messaging service, has significantly evolved since its inception in 2011¹⁶³. The app's feature set has expanded and transformed over time to meet user needs and adapt to changes in the social media landscape. In its early years, Snapchat was known for its ephemeral messaging feature, which allowed users to send photos and videos that would disappear after being viewed. This was a unique offering compared to other social media platforms at the time.

Over time, Snapchat evolved new features to enhance user experience: In 2013, Snapchat rolled out a change that allowed users to replay one snap a day, sacrificing some of its ephemerality. Geo filters, Filters, and Stickers were game changers for Snapchat. Filters and stickers allowed users to add creative elements to their snaps, while geo filters provided a location-based component, enabling users to share where they were when they took their snaps. Snapchat Discover marked Snapchat's transition from a purely social media platform to a mainstream media outlet. Launched in 2015, Discover gave users access to short-form content from major publishers, thereby expanding Snapchat's use case. Snapchat introduced the concept of 'Stories', an ephemeral feed of videos that could be broadcast to friends. They also pioneered face filters, transforming user selfies with fun and sometimes bizarre effects. More recently, Snapchat has been focusing on societal impact initiatives. For instance, during the 2020 US Election, Snapchat aimed to increase voter participation and ensure its users were well-informed through in-app initiatives.

known for more than just social media. By expanding into the metaverse and investing in technologies beyond traditional social networking, Meta aims to position itself as a leader in the evolving digital landscape.

This strategy indicates a broader focus on immersive experiences, connectivity, and the development of innovative products and services. Meta has a strong research department focused on artificial intelligence. The company is known for its calibre in AI research, which is considered competitive with industry peers like Google. Meta's commitment to AI research indicates its intention to leverage AI technologies and advancements to enhance user experiences and drive innovation within its products and services. The firm still focuses on ad revenues, however. Meta offers bid strategies and tools to help businesses achieve specific measurable outcomes, such as increasing sales, acquiring more customers, or expanding brand reach. Advertisers can choose bid strategies aligned with their primary key performance indicators to optimize campaigns and measure success based on their specific goals. It's important to note that Meta's strategy is dynamic and subject to evolution as the company continues to explore opportunities in emerging technologies, user preferences, and market demands.

The core of the DCV centres around firms being able to change. This contrasts with the RBV's focus on ensuring the firm can continue to exploit its differences to other firms over time. The common aspect of both strategic approaches is the need to be hard to copy. In the DCV hard to copy is correlated with the ability for the firm to change to continue to fit with its market better than rivals. This is a non-trivial advantage because markets change over time. In the RBV hard to copy is correlated with controlling rare resources that confer a competitive advantage. These advantages are sustainable if rivals have to pay more to acquire them, or if it is unclear how rivals can acquire them.

The next section outlines how competitive advantages based on AI can be sustainable. Importantly both the DCV and the RBV are relevant. The DCV suggests being able to evolve the firm to take advantage of AI faster, more effectively and more efficiently than rivals is a critical competency. The RBV suggests that AI deployments need to be founded in such a way that they are hard for rivals to imitate.

The hard to imitate requirement is critical for firms intending to leverage the power of AI. To date there has not been a lot of strategy in how to exploit either last gen or next gen AI developments. The commercial landscape is approaching this like it is a gold rush. That means firms are assuming first mover advantages are the most critical to developing hard to copy advantages. However resolution of complexity is a potentially better strategic approach.

• Sustaining AI Advantages

Flickr, Snapchat, and Meta are examples of low, medium, and high-water marks respectively for AI based CX. But even for Meta it is hard to see how they are going to sustain a market advantage from the AI aspects of their business models. Meta may be able to build proprietary AI tools. The question is whether other firms copy these feature sets. Companies need to think differently to exploit AI for sustainable market advantages in the context of so much AI being open source. Instead of solely relying on off-the-shelf AI deployments, or easy to copy applications, organizations should focus on data differentiation, customization and integration and domain expertise.

Data Differentiation

Building proprietary datasets that are uniquely valuable and difficult to replicate can offer a competitive advantage. Companies should collect and curate high-quality data specific to their industry or domain, enabling the training of AI models that outperform generic solutions. In some market regulators are trying to prevent this with initiatives like GDPR and the open banking initiative.

The open banking initiative refers to a practice in the banking industry that allows third-party financial service providers to access consumer banking data and information from banks and non-bank financial institutions through application programming interfaces (APIs). It involves the sharing of financial data between banks and third-party service providers, such as fintech apps, in a secure and controlled manner.

Under open banking, customers are typically required to grant consent to banks to allow access and control of their personal and financial data by third-party service providers. The purpose of open banking is to promote innovation, competition, and customer-centric services in the financial industry by enabling secure data sharing between different financial institutions.

The data a bank does not have to share can offer a competitive advantage. Non-financial data such as personal preferences, browsing history, and social media activities are generally not part of the open banking data sharing framework. The initiative typically revolves around financial transaction data and relevant customer information necessary for financial services. This data is often the most relevant for discovering customer insights and improving CX. Sadly most banks seem unable to use this type of data to customise their client interactions because they are focused on consistent scale operations. AI has the potential to break this wide open and not just for banking.

Both the manufacturing industry and energy utilities can benefit significantly from AI and proprietary data due to increased data accessibility. AI can help

optimize production processes, improve supply chain management, enhance quality control, reduce energy consumption, spread energy consumption, and enable predictive maintenance in manufacturing operations or downstream client production lines. The key here is having access to a large scale fleet of machines installed and being able to use the Internet of Things to capture proprietary performance data. A special case of this applies to transport.

AI and telemetry systems have the potential to benefit both train drivers and truck drivers by providing individual recommendations for more fuel-efficient driving¹⁶⁵. In the case of train drivers, telemetry systems collect real-time data from various sensors on the train, monitoring factors such as speed, acceleration, braking, and energy consumption¹⁶⁶. AI algorithms can then analyse this data to identify patterns and inefficiencies, generating personalized recommendations to optimize fuel usage and reduce emissions. Additionally, driver feedback systems can alert train drivers to excessive energy consumption or suboptimal driving conditions, enabling immediate adjustments. For truck drivers, telematics systems utilize onboard sensors, GPS, and connectivity to collect data on fuel consumption, engine performance, vehicle speed, and driving behaviour. AI-based analytics can analyse this data to assess driving patterns and fuel efficiency, offering personalized recommendations such as speed control, acceleration, braking techniques, and route optimization. Moreover, AI-powered systems can provide real-time feedback and coaching to truck drivers, alerting them to inefficient behaviours and suggesting improvements like reducing idling time, maintaining consistent speeds, and optimizing gear shifting for better fuel efficiency. It's not hard to imagine these systems also being deployed for ride comfort or race driving applications as well.

Companies in the e-commerce sector can leverage AI and proprietary data to improve operations and enhance customer experiences. AI can be used to personalize product recommendations, optimize pricing strategies, and streamline logistics and inventory management. The core data advantage comes from both the number of users and their spending volumes. Amazon holds a significant proprietary data advantage over other online retailers due to the higher purchase volumes of its Prime subscribers compared to non-subscriber customers. With over 200 million Prime subscribers worldwide, Amazon has access to a vast amount of customer data generated through their purchases, browsing behaviour, and preferences¹⁶⁷. More than 100 million Prime members spent an average of \$1,400 per year on Amazon in 2018, while regular customers only spent \$600¹⁶⁸ ¹⁶⁹. About two thirds of the company's customers pay for Prime. Amazon had revenues of around \$73 billion for online purchases in 2018. In comparison eBay had revenues of \$10.7B for 2018 across 170 million customers¹⁷⁰. This suggests Amazon's customers spend on average just over \$1100 per year and eBay customers

only spend just over \$70. It's not hard to see Amazon should know a lot more about its customers than eBay, even though the total number of customers is comparable. The Amazon example shows how even if the base AI platform is open source and cheap to deploy, getting value out of it may still be a scale play when it comes to CX applications designed to drive revenue. Scale is often at play in the healthcare industry with different implications.

Amazon's scale effectively confers a first mover advantage for access to differential data. The other case examples above (non-financial banking data, installed equipment performance data, telemetry data and client energy usage patterns) only require a minimum level of scale to be useful. Companies with these differential data sets can use AI to find insights to resolve complexity and potentially create sustainable advantages. This is not the case in other industries like healthcare.

The use of AI in healthcare has gained significant attention and investment in recent years, with various stakeholders, including healthcare companies, exploring its potential for competitive advantage. However, when it comes to data in healthcare, there are two scenarios that can impact the ability of healthcare companies to leverage AI for competitive advantage: data retained privately by patients and data made open source for all players.

Firstly, data retained privately by patients can pose challenges for healthcare companies looking to use AI for competitive advantage. Patients have the right to control their personal health data, and privacy regulations, such as data protection laws, ensure the confidentiality of patient information. As a result, healthcare companies may face limitations in accessing comprehensive datasets that can drive AI-powered insights and innovation. Patient reluctance to share their data or concerns about data breaches and misuse can restrict the availability and quality of data needed to develop robust AI models.

Secondly, the trend of making healthcare data open source can also impact the competitive advantage of healthcare companies. Open source initiatives aim to promote collaboration and information sharing for the collective benefit of healthcare research and innovation. While this can foster advancements in AI applications, it also means that the data used to train AI models becomes accessible to multiple players, including competitors. This openness can level the playing field and reduce the exclusivity of data-driven insights, potentially diminishing the competitive advantage that individual companies can gain from their proprietary datasets.

An example relates to skin cancer detection. Researchers conducted a study comparing the accuracy of artificial intelligence to dermatologists in identifying skin cancer. They made a database of nearly 130,000 skin disease images and trained their algorithm to visually diagnose potential cancer¹⁷¹. The AI utilized a deep learning convolutional neural network (CNN) and a

group of 58 dermatologists from around the world. To conduct the study, a test set of one hundred dermoscopic images containing melanomas and benign nevi was assembled. Initially, dermatologists were provided with the images alone, and the CNN correctly identified 95 percent of melanomas compared to the dermatologists' average of 86.6 percent¹⁷². In subsequent stages, where additional information was given, dermatologists' accuracy improved, but the CNN still outperformed them. This algorithm and the image database are open source. This means firms can use the findings for free. What they can't do is sustain an advantage with this knowledge. One way to build competitive advantage on top of open source is to customise.

Customisation and Integration

While open source frameworks provide a foundation, customization and integration with other systems can create unique solutions. Developing AI models that seamlessly integrate with existing workflows, processes, and data infrastructure can enhance efficiency and deliver specialized outcomes.

Historically Grammarly has been a writing correction tool with more capability than standard spell checkers. The latest version of Grammarly, known as GrammarlyGO, incorporates AI technology to enhance its writing features and capabilities¹⁷³. The generative AI (specifically built on OpenAI's GPT-3 large language models) provides users with various functionalities aimed at improving their writing. The context, preferences, and goals of the user are considered by GrammarlyGO to instantly generate high-quality drafts, outlines, replies, and revisions. By analysing the unique characteristics of the user, such as voice and style, GrammarlyGO tailors its suggestions and improvements to match their writing style, employing a personalized approach to elevate their writing and facilitate more effective communication.

One of the key features of GrammarlyGO is its ability to generate content based on short prompts or requests. Users have the option to provide a prompt or request GrammarlyGO to rewrite their existing content. In response, the system generates an entire draft or suggests improvements in terms of tone, clarity, and conciseness. This assists users in producing well-crafted pieces of writing. Additionally, GrammarlyGO serves as a writing assistant, aiding in unblocking ideas and enabling accelerated productivity for individuals and teams. This feature aims to overcome writer's block and stimulate new ideas for writing assignments, leveraging the power of AI to enhance the writing process and increase efficiency.

GrammarlyGO is the firm's attempt to extend its customer value proposition by expanding beyond correction to content generation. The litmus test will be how much users value the claimed ability to match the user's writing style. If GrammarlyGO gets that right then more users will convert from free to paid.

Domain Expertise

Combining AI capabilities with deep industry knowledge and expertise can create a powerful synergy. By understanding the nuances and challenges of their specific market, companies can tailor AI solutions to address industry-specific needs effectively, offering differentiated value.


Builder.ai is an AI-powered platform that enables the creation of tailor-made software applications¹⁷⁴. The platform combines the power of human expertise and creativity with AI to build a wide range of applications for diverse customers. Builder.ai allows businesses and individuals to order customized apps, similar to ordering pizza, and then uses an assembly line-like process to produce the app according to the specific requirements. The platform leverages AI to streamline the development process, helping build projects quicker and more cost-effectively.

Builder.ai is a low-code or no-code software solution, allowing users to build and manage apps without extensive coding knowledge. It aims to empower individuals, irrespective of their technical background, to unlock their potential and create applications tailored to their needs. With Builder.ai, users can choose from a template and customize the features they require, while AI calculates the price and provides fixed costs and clear timelines.

In contrast to Builder.ai's use of AI to assist customers, Amgen (a prominent biopharma player) is using AI in its quality assurance processes for visual inspection¹⁷⁵. Amgen aims to improve the inspection of various pharmaceutical products, such as syringes. Amgen partnered with Syntegon Technology (formerly Bosch Packaging Technology) to develop a fully validated visual inspection system using artificial intelligence¹⁷⁶. This project involved retrofitting an AI solution to Syntegon's syringe inspection system to advance its inspection capabilities. The AI helps identify potential defects or anomalies in products, ensuring the quality and integrity of the pharmaceuticals. AI allows Amgen to automate and optimize the inspection process, improving efficiency, accuracy, and speed. This enables Amgen to maintain high standards of quality control while reducing the time and resources required for manual inspection. The innovation was only possible by combining Amgen and Syntegon's domain knowledge to train the AI added on to the base visual inspection system.

This chapter has focused on the strategic problem of creating hard to copy, scalable advantages on top of open source AI platforms. Firms need to exploit differential data sets, customise, and integrate their offers to customers to resolve complexity and/or leverage rare domain expertise to create value market rivals cannot compete away on price. The next chapter goes beyond using AI to achieve strategic competitive advantages. It covers how customers and employees come to trust or mistrust how firms deploy AI.

VIII. Trust in AI is Complex



AI felt very strongly about keeping data private.

The literature in social psychology and marketing defines trust as the perceived credibility and benevolence of a counterpart. Trust is a key consideration for companies who want to leverage AI for several reasons. It is necessary to ensure user adoption and engagement with AI systems. If users do not trust an AI system, they are unlikely to use it, which can limit the potential benefits of the system for the company. These users are both employees and customers. There are also potential brand impacts.

Trust is necessary to protect the company's reputation and avoid negative publicity. If an AI deployment is perceived as untrustworthy, this can damage the company's brand and lead to a loss of customers and revenue. There are many examples of companies who have risked losing the trust of their customers because of their AI activities. These trust issues are relevant even when the AI deployment affects employees rather than customers.

In 2018, it was revealed that Amazon had developed an AI recruiting tool that was designed to analyse job applications and identify top candidates based on patterns in resumes and other data. However, it was later discovered that the tool was biased against women, as it had been trained on historical hiring data that was predominantly male. This led to the tool penalizing resumes that contained terms commonly used by women, such as "women's," "diversity," and "feminist." This raised concerns about the potential for AI to perpetuate and amplify biases. Amazon publicly said the tool "was never used by Amazon recruiters to evaluate candidates." The company did not elaborate further. It did not dispute recruiters looked at the recommendations generated by the recruiting engine. In this case it seemed Amazon avoided an unintended problem before they committed an error. Presumably this was to protect both the employee and consumer brand Amazon enjoys. Facebook seems to be less trustworthy in comparison.

Privacy

Facebook has faced significant criticism and scrutiny over its handling of user data. Facebook's use of AI for targeted advertising has raised concerns about the exploitation of user data and the potential for manipulative advertising practices. This has eroded trust among users who are concerned about the privacy implications of using the platform. Furthermore, Facebook's involvement in several high-profile controversies, such as the Cambridge Analytica scandal, has further damaged its reputation and eroded trust among users¹⁷⁷.

The Cambridge Analytica scandal was a data breach that involved the unauthorized collection and use of personal data from millions of Facebook users. Cambridge Analytica (a political consulting firm) used a third-party application to collect data from Facebook users and their friends without their consent. The data was then used to build psychographic profiles of users.

These profiles were used to influence political campaigns, including the 2016 US presidential election. The scandal came to light in 2018 when a whistleblower, Christopher Wylie, revealed the extent of Cambridge Analytica's data collection and usage practices.

The scandal resulted in significant backlash against Facebook and raised concerns about the company's handling of user data. Facebook was criticized for failing to protect user data and for allowing third-party developers to collect data without user consent. The scandal led to several investigations, including a US Federal Trade Commission investigation, which resulted in a \$5 billion settlement with Facebook. Overall, the Cambridge Analytica scandal highlighted the importance of data privacy and the potential for misuse of personal data in the age of big data and AI. It also implied a need for greater regulation and oversight of companies that collect and use personal data. Despite Facebook's efforts to address these issues, the platform still faces significant challenges in rebuilding trust with its users. One issue relates to the potential for organisations to manipulate groups of people at scale when they have consented to give their data to a platform like Facebook. Another issue relates to lack of consent.

Clearview AI, an American facial recognition company, was fined €20 million by the French data protection authority (CNIL) for illegally collecting and processing biometric data belonging to French citizens¹⁷⁸. Clearview AI collected publicly available images and videos of people from websites and social media platforms and associated them with identities, thereby creating a biometric database of facial scans and identities. The company sold database access to various operators of facial recognition systems. General Data Protection Regulation (GDPR) dictates that any data collection must be communicated to people clearly and requires consent. Clearview's model meant individuals were unaware their images were potentially being used for identification by Clearview AI clients. The trust issue in this case seems obvious but it is not necessarily because of the GDPR violation. The case does relate to the lack of informed consent. But it also relates to payment.

Clearview collected two sources of publicly available data and connected them together to create a database with valuable information. It sold this database to others. What it did not do is pay the people who were identified in the database. Ultimately this was deemed a violation of the GDPR regulations by the French authorities. This suggests the motivation for GDPR being enacted is not just related to privacy but also about who profits from data collection and usage.

The violation relating to payment in this case seems inconsistent with the way that journalists have been allowed to operate historically. (Disclaimer: the author is not qualified as a lawyer in this area, so perhaps naively or foolishly

will try to apply some common sense to this situation to assert precedent). Historically journalists and photographers, including paparazzi have been allowed to photograph people in public places and sell their pictures to news organisations without the consent of, or any payment to, the subjects in their photographs. They could also link the photos to information accessed from subjects' social media sites. They could do this even if they were only freelancers. It seems Clearview simply did a similar thing at scale but is in violation of the GDPR laws.

GDPR has been criticised because it does not work as it was billed to. Here is an excerpt from a 2019 article:

The GDPR was billed as the gold standard of data protection, offering the strongest data rights in the world. It has forced companies everywhere to modify their operating models, often at great cost...

...Yet for those of us living under the GDPR, what has really changed?

... The only difference is that now we are forced to participate in our own privacy violation in a grotesque game of "consent".

Most websites nudge us into clicking "I consent" by making it harder for us not to. Those that do offer an "I do not consent" option force us to navigate a complicated menu of privacy settings, all of which offer only the veneer of privacy...

... they are betting that most of us will choose convenience over data protection. And so we click "I consent" to cookies and other web trackers that follow us around, creating an ever-growing digital self that is monitored, used, bought and sold...

...The GDPR was supposed to prevent all of this...

...Nor is the GDPR stopping the construction of a surveillance society – in fact, it may even legalise it. The collection of biometric data, which occurs with facial recognition technology, is prohibited under the GDPR unless citizens give their explicit consent. Yet there are exceptions when it is in the public interest, such as fighting crime.

This is how an exception becomes the rule...

GDPR has not stopped consumers from being tracked nor has it resulted in them generally being paid for their data.

Effectiveness

The fallibility of AI also brings into question our ability to trust it. NEC's facial recognition technology, NeoFace, has been developed to identify individuals in real-time, even in crowded environments¹⁷⁹. NeoFace has won multiple awards for performance. NEC claims high accuracy of NeoFace's image matches even with poor-quality images from compressed surveillance video. UK Police forces have performed trials in public places such as sports stadiums, shopping malls, and crowded streets. Ten trials with NeoFace have been completed to date using dedicated camera equipment. The tech has not performed well. In eight Police trials 2016 and 2018, 96% of facial recognition matches misidentified innocent bystanders as watch list subjects. In one instance filmed by the BBC, a passer-by in North London covered his face, yet police still photographed him and fined him for disorderly conduct. Other trials conducted by the South Wales Police were also problematic because 91% of matches misidentified innocent people.

The trials led to legal action taken against South Wales Police by a member of the public who claimed that the technology violated their right to privacy, interfered with their right to protest, and breached data-protection laws. In addition to concerns over privacy and data protection, the use of facial recognition technology raises questions about accuracy and bias. NeoFace has been found to be less accurate when identifying people with darker skin tones.

In Australia Kmart and Bunnings have suspended the use of facial recognition technology in their stores due to an ongoing investigation by Australia's privacy regulator, the Office of the Australian Information Commissioner (OAIC)¹⁸⁰. Last month, Consumer group Choice revealed the two companies were using the technology to capture images of people's faces from video cameras as a unique faceprint to protect customers, staff and reduce theft in select stores. The companies are now being investigated over their use of the technology and whether it is consistent with privacy laws. Bunnings' managing director has confirmed that the company has ceased using the technology, which was only used to identify when a person who had been banned from Bunnings stores entered the store. Kmart has also temporarily ceased using the technology in its small number of trial stores.

To address these concerns, some jurisdictions have implemented regulations that require companies to obtain explicit consent from individuals before using their images to train facial recognition algorithms. In addition, some researchers have developed methods for testing facial recognition technology on synthetic data or on images that have been altered to protect the privacy of the individuals in the images.

Transparency

Going beyond consent, privacy and effectiveness issues leads to other cases where AI has been used in a way that corrodes trust because of a lack of transparency. In 2019, Volkswagen (VW) was accused of using AI to create emissions cheating software. The AI detected when the car engine was being tested for emissions and changed to more efficient, less polluting mode to get better results on the test. VW management asked engineers to develop and install the software into their diesel vehicles, because these cars could not pass US emissions tests without them. VW deliberately sought to conceal this software. The company was heavily criticized for its actions, which ultimately led to a \$30 billion settlement with various governments and consumers. The emissions cheating scandal demonstrated how AI can be used to deceive regulators and the public, leading to significant harm to both the environment and public health. The brand damage to VW was considerable.

Implications

Sometimes the trust issue relates to the potential implications of a legal use of an effective technology with proper consent, privacy protections and transparency. Google has used AI legally but not always in a way supporting the firm's 'do no evil' aspiration. In 2018, the search engine giant came under fire for its involvement in the US Department of Defense's Project Maven, which used AI to analyse drone footage¹⁸¹. Google employees and external groups expressed concern that the company's involvement in the project would contribute to the development of autonomous weapons, which could have severe ethical and human rights implications. The controversy ultimately led to a decision by Google to withdraw from the project and commit to developing AI principles that prioritize ethical considerations. Legally Google had every right to continue the work. The potential impact on reputation was another matter.

Google stated then it would no longer develop AI for use in weapons but would continue to pursue defence contracts in other areas. Since then, the company has established a significant line of business including contracts for detecting corrosion on Navy vessels using machine learning, supporting aircraft maintenance for the Air Force, and providing cloud security technology to the Pentagon's Defense Innovation Unit. In 2020, Google won a portion of a large CIA cloud contract, and in 2021, it jointly won a \$1.3 billion deal with Amazon to supply cloud services to the Israeli government, including its defence agencies. While these contracts have not drawn the same level of scrutiny as Google's involvement in Project Maven, they illustrate the company's continued involvement with the defence industry. It seems it is ok to continue with the Google Self-Driving Car Project with its potential to displace employee drivers because this is different to creating

machines designed to remove some armed forces personnel from situations they could be injured or killed at the expense of enemy forces¹⁸². In all these projects Google has acted legally to develop effective technology. They operated transparently with risk of privacy violations or a lack of informed consent. The question is whether these projects are ethically ok.

It's not enough to assume regulatory compliance equates to trust. Compliance is necessary but not sufficient. Many countries have regulations around data privacy and ethical AI use, and failure to comply with these regulations can result in significant legal and financial consequences. Customers are quite aware legally compliant does not automatically mean trustworthy. Even where legal violations do not occur and there are no issues of informed consent and privacy and data protections, there are still other issues to consider. AI tech that is not perfect creates potential trust issues. For example, the roll out of self-driving cars has been limited because the fear is they are not perfectly able to avoid accidents. It doesn't matter that they can probably drive a lot better than many people already. Even if they could pass a driving test to get a licence many people would still not be happy to trust such vehicles on the road. Maybe one day.

The flip side of this issue arises when autonomous vehicles are perceived to be safe enough to allow them on public roads. This eliminates many jobs. Already such vehicles are being used off road in various industries, including mining, agriculture, and transportation. In the mining industry driverless trucks are being used in operations to transport materials such as coal, iron ore, and copper by firms like Rio Tinto and BHP. In agriculture driverless tractors are being used to perform tasks such as planting, seeding, and harvesting. John Deere, one of the world's largest farm equipment manufacturers, has been developing autonomous tractors for several years. Driverless forklifts and other autonomous vehicles are being used in warehouses and logistics centres to move goods and materials. Amazon is one company that has been using autonomous robots in its warehouses for several years. Driverless taxis and delivery trucks are currently being tested for transport use cases. Waymo and Tesla are developing and testing autonomous vehicles for public transportation. Driverless vehicles are becoming increasingly common in many industries, and their use is likely to continue to grow in the coming years as the technology continues to improve.

• Building Trust

Many researchers have examined the concept of trust, what it means and how it is established. In 2021 the University of Queensland and KPMG published research *Trust in Artificial Intelligence - A Five Country Study*¹⁸³. The report found trust is a crucial factor for the acceptance and widespread use of AI, and there are four key drivers that influence citizens' trust in AI systems,

namely, beliefs about the adequacy of current regulations and laws to make AI use safe, the perceived impact of AI on jobs, familiarity and understanding of AI, and the perceived uncertain impact of AI on society. The perceived adequacy of current regulations and laws is clearly the strongest of these drivers, highlighting the importance of ensuring adequate regulatory and legal mechanisms are in place to protect people from the risks associated with AI use. Despite low levels of trust in AI systems, citizens generally accept or tolerate AI, but few approve or embrace it, and some outright reject AI. Citizens' trust and support of AI depend on the purpose of the AI system, with the public being more trusting and supportive of AI use in healthcare than in human resources.

Most citizens expect AI to be regulated and to have external, independent oversight. Most citizens across the five countries disagree or are ambivalent that current regulations and laws are sufficient to make the use of AI safe. Therefore, there is a need to strengthen and communicate the regulatory and legal framework governing AI across all surveyed countries. Citizens in each country have very clear expectations of the principles and related practices organizations deploying AI systems should uphold to be trusted, such as performance and accuracy, data privacy, security and governance, transparency and explainability, accountability, risk and impact mitigation, fairness, and human oversight.

Most citizens would be more willing to use AI systems if assurance mechanisms were in place, such as independent AI ethics reviews, AI ethics certifications, national standards for transparency, and AI codes of conduct. Furthermore, citizens want to know more about AI but currently have low awareness and understanding of AI and its uses. The survey results suggest a strong need and appetite for a public AI literacy program.

Collectively, these survey insights provide an evidence-based pathway for building and maintaining the trust and acceptance of AI systems by citizens of western nations. This pathway requires government and business to act by living up to citizens' expectations of trustworthy AI, strengthening the regulatory framework for governing AI, and enhancing AI literacy of the public and employees. The survey insights are relevant for informing AI policy and practice within business, government, and NGOs at the national level, as well as multinational and pan-governmental AI policy and practice.

The case studies presented earlier in this chapter suggest regulations alone will be inadequate for building trust in AI. The following sections introduce trust in business and technology as foundations for considering trust in AI.

Business Trust

One of the most well-known models for establishing trust in commercial transactions is the “trust transfer model” proposed by Doney and Cannon (1997) in their seminal paper "An Examination of the Nature of Trust in Buyer-Seller Relationships". The research seems generalisable to C2C contexts even though the research sample was limited to B2B transactions. The model suggests that trust can be transferred from a trusted through a process of trust-building activities, such as reputation, information sharing, and structural assurances. In this model, trust is seen as a multidimensional construct that includes competence, benevolence, and integrity. At the core of trust is the idea that one side is at risk and needs to rely on the other party to remain safe or unexploited. This is where competence, benevolence and integrity apply. Trust starts with the idea that the trusted party can do what is required. Benevolence means the trusted party will act in the interests of the trustor. Integrity relates to communications where statements of competence and commitments to benevolence are made. The trusted party must remain true to its word and do what is required. All three dimensions are essential for establishing trust in a commercial transaction.

Technology Trust

In addition to the trust transfer model, other models have been proposed for establishing trust in e-commerce, such as the technology acceptance model (TAM) and the unified theory of acceptance and use of technology (UTAUT). These models focus on the role of technology and its influence on trust-building activities.

The Technology Acceptance Model (TAM) is a theoretical framework used to explain how and why individuals adopt and use technology¹⁸⁴. The model was first proposed by Davis in 1989 and has since been widely used in the field of information systems research. The model proposes that an individual's intention to use a technology is influenced by two main factors: perceived usefulness and perceived ease of use. Perceived usefulness is the extent to which an individual believes that a technology will improve their performance or productivity, while perceived ease of use is the extent to which an individual believes that a technology is easy to use and understand.

Several studies have confirmed the validity and reliability of the TAM model across a variety of settings and technologies. For example, a study by Venkatesh and Davis (2000) found that perceived usefulness and perceived ease of use were significant predictors of individuals' intentions to use e-commerce technologies¹⁸⁵. Another study by Wu and Wang (2005) found that perceived usefulness and perceived ease of use were significant predictors of individuals' intentions to use mobile banking services¹⁸⁶.

Overall, the TAM model provides a useful framework for understanding how individuals perceive and adopt new technologies. By identifying the factors that influence technology adoption, researchers and practitioners can develop strategies to promote technology acceptance and use.

The Unified Theory of Acceptance and Use of Technology (UTAUT) model proposes that user acceptance of technology is influenced by four key factors: performance expectancy, effort expectancy, social influence, and facilitating conditions¹⁸⁷. Performance expectancy refers to the user's belief that using the technology will improve their performance, while effort expectancy refers to the perceived ease of use. Social influence refers to the user's perception of the opinions of others, and facilitating conditions refer to the user's belief in the availability of resources and support for using the technology. The UTAUT model also suggests that these factors can be influenced by various individual and contextual factors.

The UTAUT model has been extensively studied in various contexts, including healthcare, education, and business. For example, a study by Al-Qeisi et al. (2019) applied the UTAUT model to investigate the factors influencing the adoption of mobile banking in Jordan¹⁸⁸. The study found that performance expectancy, effort expectancy, and facilitating conditions were significant predictors of user intention to adopt mobile banking. Another study by Fathema et al. (2015) applied the UTAUT model to explore factors affecting the acceptance of mobile learning among university students in Bangladesh. The study found that social influence, facilitating conditions, and effort expectancy were significant predictors of students' intention to use mobile learning. The UTAUT model has proven to be a useful framework for understanding user acceptance and adoption of technology in various contexts, providing insights for the design and implementation of technology-based interventions.

• Trusting AI

The mechanisms for building trust between a company and its customers are like those for building trust in AI. Both require transparency, reliability, accountability, and ethical considerations. Companies need to be transparent about their practices and policies and communicate clearly with customers to build trust. Similarly, AI systems need to provide transparent and understandable explanations for their decisions to build trust with users.

However, there are also some differences between building trust in a company and building trust in AI. One difference is that building trust in a company often involves building a brand and reputation over time, whereas building trust in AI systems requires demonstrating trustworthiness through technical design and functionality. Additionally, building trust in a company often involves building personal relationships with customers and establishing

a sense of community, whereas building trust in AI systems is more focused on establishing credibility and reliability.

In both cases, building trust requires a focus on the user or customer experience. Companies and AI designers need to consider the needs and expectations of their users and customers to build trust and establish a long-term relationship. This requires a willingness to listen to feedback and adapt to changing circumstances, as well as a commitment to maintaining high ethical standards and prioritizing the user's best interests.

Building trust in AI is different from building trust in other areas because AI systems are complex and often opaque. Users may not understand how an AI system works, how it makes decisions, or how their personal data is being used. This lack of understanding can lead to suspicion and distrust. Personal data is a key factor in building trust in AI. Users need to feel confident that their personal data is being used ethically and securely. This requires transparency around data collection, use, and storage, as well as clear communication about the purpose of data collection and how it benefits the user. Companies and AI designers also need to ensure that users have control over their personal data and that it is not being exploited for commercial gain.

Value exploitation is another factor that can erode trust in AI. Users may feel that AI systems are exploiting their personal data or using it to manipulate their behaviour. Companies and AI designers need to be transparent about the value exchange that is taking place and ensure that users feel that they are getting a fair deal. This requires clear communication about how the AI system benefits the user and how this compares to the benefits the company is receiving.

Understanding is a crucial factor in building trust in AI. Users need to be able to understand how the AI system works, how it makes decisions, and what the implications of those decisions are. This requires clear and understandable overview of the AI system's algorithms and decision-making processes, as well as user-friendly interfaces that make it easy for users to interact with the system.

Finally, use cases are an important consideration when building trust in AI. Different use cases require different levels of trust and different approaches to building trust. For example, an AI system used in healthcare may require a higher level of trust than an AI system used for entertainment. Companies and AI designers need to be sensitive to the specific needs and expectations of users in each use case and tailor their design and communication strategies accordingly.

• AI Trust Cohorts

It is useful to classify people based into four groups based on their level of trust in AI. The first group are those who mistrust AI because they believe its use will disadvantage them. This group can be divided into two sub-groups: customers and employees. Customers in this group are concerned that their personal data will be misused, and that AI may be used to discriminate against them. They may believe AI will create price differences based on their observed behaviour and set higher prices for certain groups. Employees, on the other hand, are worried about job displacement and the potential for AI to replace them in their current roles. They may also be concerned about the use of AI for decision-making in the workplace, such as in hiring and promotion decisions. AI deployments may be at risk of bias against certain groups, leading to unfair outcomes. Approximately 22% of citizens believe that AI will create more jobs than it eliminates. Nevertheless, about 45% of employed citizens indicate that they use AI in their work, although they use it infrequently. Most citizens (ranging from 70% to 76%) are at ease with using AI at work for task automation and augmentation. However, they are less comfortable with using AI for employee-focused tasks, such as monitoring, evaluating employees, recruitment, and selection.

The second cohort mistrusts AI because they don't understand how it works and/or don't believe it works better than the alternatives. This lack of trust can be attributed to the liability of newness, which refers to the inherent scepticism that people have towards new technologies until they become more widely adopted and better understood.

The third cohort is unable to trust AI simply because they are unaware of its deployment. This may be due to a lack of awareness or education around AI or a lack of transparency from organizations using AI. This cohort may not have had any exposure to AI or may not understand how it is being used in their daily lives.

The final group trusts AI because they believe its use advantages them. This group sees AI as a tool that can improve their lives, whether through better healthcare outcomes, increased efficiency, or more personalized services. They may be early adopters of new technologies and may be more willing to take risks and experiment with new AI applications.

Classifying people into different groups based on their level of trust in AI can be useful for understanding how to encourage people to adopt AI. By understanding the reasons why people may be hesitant to use AI, it is possible to design strategies to address their concerns and increase their confidence in the technology. Here are some strategies that organizations can use to gain AI adoption for each cohort based on their attitudes towards AI:

Cohort 1 (mistrust AI because they believe its use will disadvantage them): For customers in this group, organizations can provide clear and transparent information on how their personal data will be used and protected. They can also provide assurances that AI will not be used for discriminatory purposes. For employees in this group, organizations can provide training and education on how AI can augment their jobs and how they can work alongside AI to improve productivity and efficiency. Organizations can also provide job training and reskilling programs to help employees transition into new roles if necessary. To address the bias issue, organizations can conduct regular audits of AI systems and processes to identify and mitigate bias and ensure fair outcomes.

Cohort 2 (mistrust AI because they don't understand how it works and/or don't believe it works better than alternatives): Organizations can provide training and education on AI and its capabilities. They can also offer demos or trials of AI systems to show how they work and how they can benefit users. Organizations can provide case studies and success stories of other companies or individuals who have successfully implemented AI in their work or personal lives. This can help build trust and confidence in AI.

Cohort 3 (unable to trust AI simply because they are unaware of its deployment): Organizations can increase transparency around AI deployment and provide more information on how it is being used in their daily lives. This can be done through educational campaigns, social media, and other communication channels. Organizations can also collaborate with policymakers and regulators to develop guidelines and regulations around AI deployment and use. This can help increase transparency and build trust in AI.

Cohort 4 (trust AI because they believe its use advantages them): Organizations can highlight the benefits of AI and its potential to improve their lives. This can be done through targeted marketing campaigns and messaging that focus on the specific benefits that AI can provide. Organizations can also offer incentives for early adopters of AI, such as discounts or exclusive access to new AI applications or services. This can help encourage users to try AI and see its benefits for themselves.

- **Increasing AI Trust**

One example of where AI deployment has increased customer trust is in the financial industry, particularly in fraud detection and prevention. AI-powered fraud detection systems can quickly analyse vast amounts of customer data and detect fraudulent activity in real-time, thereby reducing the risk of fraud for customers and increasing their trust in the institution.

For instance, banks and credit card companies use AI algorithms to detect unusual transactions and flag them for review. This system can identify a

broad range of fraudulent activities such as account takeover, card skimming, and phishing scams. It can also analyse a customer's typical behaviour and flag any deviation from the norm, such as transactions in a foreign country or large purchases.

Nubank, a Brazilian fintech company, has implemented AI-based fraud detection systems to monitor all its customers' transactions in real-time¹⁸⁹. The system, called "Sereia," is powered by machine learning algorithms that can detect suspicious patterns and activities that may indicate fraudulent behaviour. The system automatically flags any transactions that appear to be suspicious, and a team of fraud analysts then reviews these cases to determine whether they are indeed fraudulent or not.

Additionally, Nubank also uses AI to assess the creditworthiness of its customers and to prevent loan defaults. The company's machine learning algorithms analyse a wide range of data points, including credit history, employment status, and income, to determine a customer's ability to repay a loan. By using AI to assess credit risk, Nubank can offer loans to a wider range of customers who may have previously been deemed too risky by traditional banks. This has helped the company to grow rapidly and become one of the largest fintech companies in South America.

One example of AI deployment increasing customer trust in Europe is in the healthcare industry. Berlin-based healthtech company Ada Health has developed an AI-powered personal health companion app that uses natural language processing and machine learning to help people better understand their symptoms and possible diagnoses¹⁹⁰. The app asks users a series of questions to identify potential health issues and provide personalized health recommendations.

By using AI to provide a more accurate and personalized healthcare experience, Ada Health has been able to increase customer trust and satisfaction. In a survey of over 1,000 users, 92% of respondents said they felt more informed about their health after using the app, and 86% said they felt more confident in their ability to manage their health¹⁹¹. Additionally, the app has received positive reviews and high ratings from users on both the Apple App Store and Google Play.

The deployment of AI in the healthcare industry, particularly with personal health companion apps like Ada Health, is becoming increasingly common in Europe and around the world. These apps have the potential to significantly improve patient outcomes and experiences, while also increasing customer trust in healthcare providers offering AI services.

- Damaging AI Trust

There have been allegations that Uber Australia has been creating fake rides in the app to deceive riders into waiting for a ride that never arrives, and then switching the ride at the last minute to keep them waiting longer. This practice has led to frustration and decreased trust among riders. Some drivers suspect that the app occasionally creates fake rides and cancels them last minute to keep drivers in specific areas until a ride is available.

Adding to this, Uber has increased fares to be more in line with traditional taxi fares, while also increasing the frequency and magnitude of surge pricing. This has led to further frustration among riders, who feel that they are being taken advantage of. Uber does not provide phone support or open email support for riders. To make a support request the rider must fit their issue into the triage platform. This requirement, the misinterpretation of some support requests and the repetitive natures of some support replies all suggest that rider support may not be manned by human customer service representatives. The user interface names its support agents to imply riders are interacting with a real human being. It is also impossible to escalate an unresolved complaint to a supervisor or change to a phone call or email channel.

Overall, the combination of these factors - the alleged creation of fake rides, increased fares and surge pricing, lack of phone support, and suspected automated chatbot support masquerading as human - has led to a significant erosion of trust in Uber Australia among both riders and drivers.

Here is a record of one suspect support interaction from earlier this year. You can see from the replies that it appears the alleged chatbot involved never really understood what the complaint was about. Here are the clues: There was no way to get to escalate the complaint to a real person. The 'Community Operations specialists' don't exist or didn't bother to make contact. This was the 'original complaint' referred to in the first support response. The last support response is the most telling because the complaint was already on the website referred to in that reply. While unlikely, if these responses really were from human support, they were no better than low quality automated software. Does this call into question Uber Australia's trustworthiness for you?

My driver or their vehicle didn't match the profile in my app

Wed, Mar 08, 5:12 PM

Did your driver match the profile photo?: true

Did you complete the trip?: true

Vehicle colour, make, and model: White car arrived different to blue car first scheduled.

Was the license plate different to the one displayed in your app?: Yes and no. 2nd car is as per revised trip. Read details below

Additional details: The app first seemed to be sending me a blue car that never arrived. I waited for 15 mins. I watched the 'car' in the app on the map arrive to be apparently right in front of me. Because it was blue and no other blue cars were on the street it would have been impossible to miss. But no car actually arrived and then the blue car with Mohammed as the driver magically disappeared in the app and a new car was allocated. And I had to wait another 15mins to get picked up. I believe the first car was never on route but the plan was to trick me into waiting longer. This is not the first time this has happened. Not happy at all. This is deceptive.

Rider. - Mar, 5:12 pm

Thank you for reaching out and for letting us know about this incident and I'm sorry to hear about this experience.

I would like to reassure you that we are treating this with the utmost seriousness. A member of our team will be in touch with you shortly.

In the meantime, if you feel additional information could be helpful, please reply to this message.

If you have already been in contact with one of our Community Operations specialists regarding a safety incident and feel that you have additional information that could be helpful, please respond to our original message and we will continue to help you there.

Lavneet - Mar, 5:14 pm

L

I was very angry when this happened and could not find a way to make this complaint. I have since worked out how to make the complaint again in the correct tab before seeing your reply here. Please note my ultimate driver Nirav did nothing wrong.

The problem is with Uber. I think this situation was created to get me to wait for 30 minutes by misrepresenting there was a car on the way when there was no car on the way. The reason I think this is there was no blue car on the street in front of me when the app showed the car at my location. I also didn't get a driver cancelled notification as I have had on some trips in the past. So this seems like a scam to trick me into waiting longer than I would have if I had known the real length of the delay. This is not the first time something like this has happened, but it is the first time I could see that a car apparently on the way to me did not actually exist. I say shame on your organisation for this deception. And shame on Uber for making it so hard to find how to make a complaint like this that does not fit in the support menu structure in the app. There should be at least a general email address you can make a complaint to without the triage crap this has put me through.

Rider. - Mar, 10:09 pm

Hi XXXXXX,

Sorry to hear that your driver did not arrive on time for pickup on this trip and you had to wait. Let me look into this for you.

We know that your time is valuable and that additional wait time can be frustrating when you're ready to ride, please know that we're working to improve ETA accuracy.

We've noted and appreciate your feedback as we continue working to improve the rider experience.

Your understanding in this regard is appreciated.

Gaurav - Mar, 10:11 pm

G

This is not ok. It is not ok to resolve this ticket. You have not addressed my concern. It is not about the wait time. It is about the deception. I want to actually talk to the person who is responsible for this. I want a manager to phone me.

Rider. - Mar, 10:13 pm

Thu, Mar 09, 10:13 AM

Hi XXXXXX,

We're sorry to hear about this pickup. This is definitely not the kind of experience we'd like you to have on Uber.

Sometimes cancellations are unavoidable. It is possible your driver may have a valid reason for cancelling the trip. It remains however that this pickup did not go as smoothly as expected and may have given you inconvenience.

Rest assured that this feedback has been taken into account. Thank you for taking the time to share this with us.

We hope your next ride with us goes smoothly!

Renz - Mar, 10:13 am

R

Please respond to my actual complaint. The problem was there was no driver and car the first time. Uber sent me a fake ride. This is deceptive. Please have a real person call me to discuss.

Rider. - Mar, 10:25 am

Hi XXXXXX,

Thanks for letting us know about this issue. We know it can be frustrating when there are pickup difficulties with the driver-partner and we're sorry to hear about your experience.

The way to contact Uber support will be directly through the Uber app or from your computer at help.uber.com. We do not have a support phone number.

Once you submit your issue, you'll be able to continue the conversation through the app from the HELP menu.

Contacting us via HELP in the app or at help.uber.com allows us to collect more information so that we can resolve your issue more quickly.

We appreciate your feedback and we hope that your future experiences will be much more consistent and reliable moving forward.

Renz - Mar, 10:33 am

R

This is me already contacting you for help at help.uber.com. I am now very tired of these canned responses. I need to get a response from a human. How do I escalate my issue past you to talk to an actual person not a bot?

Rider. - Mar, 4:21 pm

Please escalate my complaint to a real human

[No further reply received at time of writing a month later]

• Critical AI Trust Issues

One critical trust warning for deploying AI is not trying to pass software off as a human being. This may be very tempting for some organisations as the AI capability to emulate humans gets better. Especially in interaction channels limited to text only exchanges. This is a special case of lack of informed consent. And it's not necessary. With quality natural language processing humans will suspend their disbelief they are interacting with software. This occurs in much the same way we can get so consumed watching a movie or reading a book that in our heads we are there with every twist and turn of the story. We know it's not real, but we go along with the experience anyway. Several humans who are interacting with ChatGPT have reported to this author that they find themselves including please in their prompts to the software! The examples of AI with some emulation of theory of mind above (including Woebot, Eliza, the NPCs in Eve Online and Replika) are effective as interactive companions even though it is obvious to their human user counterparts that they are not organic minds. AI implementations seriously breach trust when they are passed off as humans.

A second critical trust warning for deploying AI is about customer exploitation. Some companies use AI-powered chatbots to engage with customers and collect data on their preferences and behaviour. While chatbots can be a useful tool for providing customer service and support, they can also be used to manipulate customers into making purchases they may not need or cannot afford. AI can be used to analyse large amounts of data on consumer behaviour and tailor sales and marketing strategies to target specific groups of people.

A simple theoretical framework developed in a 2021 study ... can be used to assess behavioural manipulation enabled through

AI¹⁹². The study mostly deals with users' "prime vulnerability moments", which are detected by a platform's AI algorithm. Users are sent ads for products that they purchase impulsively during these moments, even if the products are of bad quality and do not increase user utility. The study found that this strategy reduces the derived benefit of the user so that the AI platform will extract more surplus, and also distorts consumption, creating additional inefficiencies.

The US chain store Target has used AI and data analytics techniques to forecast whether women are pregnant to send them hidden ads for baby products¹⁹³. Uber users have complained that they pay more for rides if their smartphone battery is low, even if officially, the level of a user's smartphone's battery does not belong to the parameters that impact Uber's pricing model. Big tech firms have often been accused of manipulation related to the ranking of search results to their own benefit, with the European Commission's Google shopping decision being one of the most popular examples.

It is not surprising to discover AI researchers are beginning to work out how AI systems can learn to manipulate humans. Dezfouli, Nock and Dayan focused on this in their research published in 2020¹⁹⁴. Here's an overview of their work with **emphasis** added by this author.

Significance

"What I cannot efficiently break, I cannot understand."

Understanding the vulnerabilities of human choice processes allows us to detect and potentially avoid adversarial attacks. We develop a general framework for creating adversaries for human decision-making. The framework is based on recent developments in deep reinforcement learning models and recurrent neural networks and can in principle be applied to any decision-making task and adversarial objective. We show the performance of the framework in three tasks involving choice, response inhibition, and social decision-making. In all of the cases the framework was successful in its adversarial attack. Furthermore, we show various ways to interpret the models to provide insights into the exploitability of human choice.

Abstract

Adversarial examples are carefully crafted input patterns that are surprisingly poorly classified by artificial and/or natural neural networks. Here we examine adversarial vulnerabilities in

*the processes responsible for learning and choice in humans. Building upon recent recurrent neural network models of choice processes, we propose **a general framework for generating adversarial opponents that can shape the choices of individuals in particular decision-making tasks toward the behavioural patterns desired by the adversary.** We show the efficacy of the framework through three experiments involving action selection, response inhibition, and social decision-making. We further investigate the strategy used by the adversary in order to gain insights into the vulnerabilities of human choice.*

The third critical trust warning for deploying AI is about the potential to create unemployment. With the rise of automation, there are concerns about the future of employment and how it will impact the global workforce. As more jobs become automated, there is potential for individuals to take on more complex roles, transitioning from physical labour to cognitive tasks.

For example, the trucking industry currently employs millions in the United States alone. However, the advent of self-driving trucks promised by Tesla's Elon Musk could displace these workers in the coming decade. While there are concerns about job loss, the increased safety benefits of self-driving trucks also make them an ethical choice.

There is ongoing debate among researchers and experts about whether AI automation creates more jobs than it removes. Some studies suggest that automation can lead to job displacement in certain industries, while others suggest that it can create new job opportunities and improve overall productivity.


One study by Gartner suggests that AI automation will create 2.3 million new jobs by 2020, while eliminating 1.8 million jobs, resulting in a net gain of 500,000 jobs¹⁹⁵. The study argues that automation will lead to new job opportunities in areas such as healthcare, education, and the public sector, and that it will also increase efficiency and productivity in many industries.

Another study by the World Economic Forum suggests that automation will lead to a shift in job roles rather than a net loss of jobs¹⁹⁶. The study argues that automation will eliminate some low-skilled jobs but will also create new opportunities in areas such as data analysis, cybersecurity, and human-machine interaction.

However, other studies suggest that the impact of automation on employment is more complex and depends on factors such as industry, geography, and skill level. Some experts argue that automation may lead to job displacement in certain industries, such as manufacturing and transportation, while creating new opportunities in others.

Managing AI trust issues will be critical for business if the real benefits from AI are to be maximised. The starting point is ensuring user data privacy and transparency are in place for effective AI deployments. After that the implications of how AI will change the commercial contexts becomes important. Passing off AI as human, using AI to exploit customers and ignoring job displacement will provoke regulatory reaction. Organisations that get past the temptation to act with a short-term focus have good reasons to undertake strategic foresight related to the future of AI. That is the subject of the next and penultimate chapter.

IX. Future Outlook:



AI just couldn't see it coming.

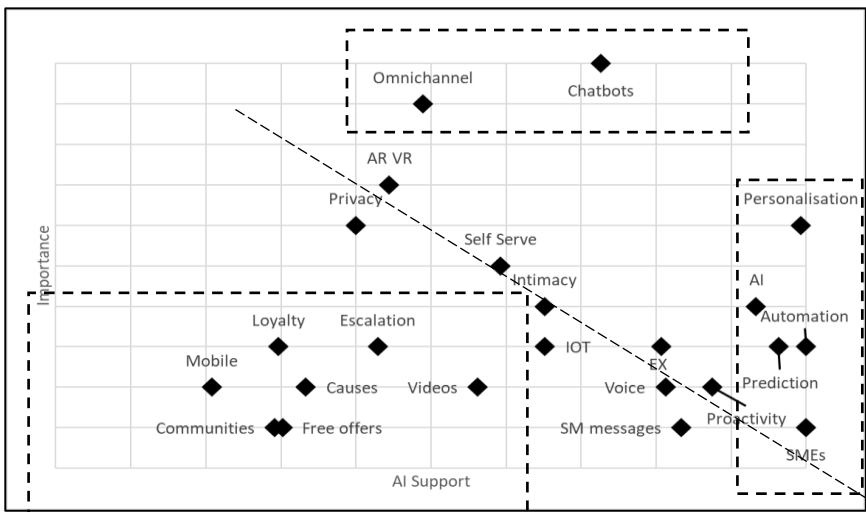
This chapter starts with an important disclaimer: it is often foolish to try to predict the future because of the uncertainties involved. The future is only influenced not determined by the past. So this chapter is an attempt to extrapolate some of the influences related to customer experience because of how AI capability is rapidly evolving. It is unlikely all the implied predictions will come to pass. Instead it is better to treat this chapter as a guide to what can be reasonably expected.

• Trend Analysis

The following trends were researched by analysing the results of two Google searches in May 2023. The first search was for “most important trends in CX”. The first 8 non-sponsored websites with useful information were analysed to identify future expectations. These were then groups into distinct cohorts to elicit 23 most significant CX trends¹⁹⁷. A count of the number of sites each CX trend appeared on was retained as a proxy for the importance of the trend. The validation for this approach is based on the idea that Google is in the business of providing the best answers to searches like this.

A similar search was conducted for “most important trends in AI” and this was used to build a construct for how much each of these trends supported the CX trends¹⁹⁸. The results are shown on the scatter plot below. The table underneath the graph contains the raw data and constructs for the graph.

FIGURE 11: MOST SIGNIFICANT CX TRENDS SUPPORTED BY AI



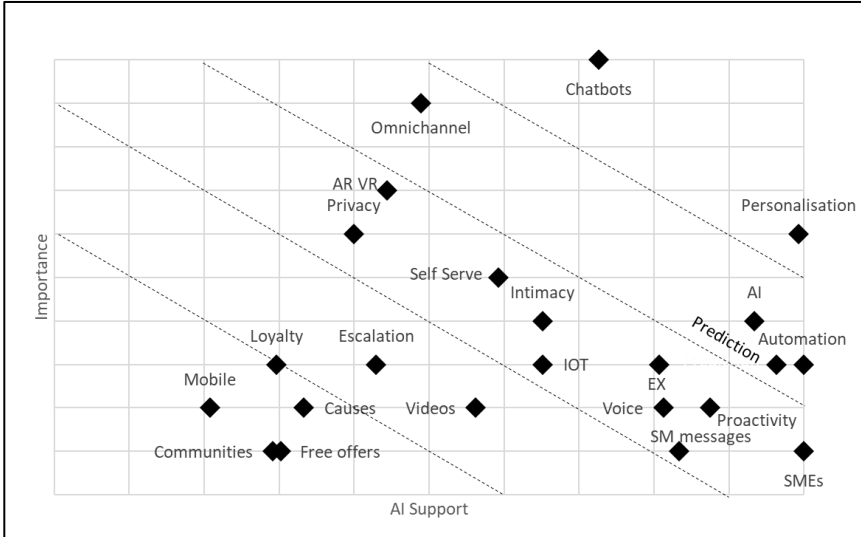
		MOST SIGNIFICANT CX TRENDS														AI Support	Normalised	
		Automation	Vertical use cases	Ethical AI and regulation	Cybersecurity	Collaborative AI	Generative AI	LLMs	Low code No code	IoT digital twins	Multimodal	Quantum computing	Personalisation AI	Predictive analytics	Transparency	24.5	7.3	
	Websites	6	6	5	5	4	4	3	3	2	2	2	1	1	1	1	13.5	4.0
Chatbots	10	1	0.25			0.5	1	1	1	0.5	1		1	1		0.5	13.5	4.0
Omnichannel	9	0.5				1	0.5	1		0.5	1		1		0.5	15	4.4	
AR VR	7					1	1	1		0.5	1		0.5			13.5	4.0	
Privacy	6		0.25	1	1											33.5	9.9	
Personalisation	6	1	0.5	1		1	1	1	1	1	1		1	0.5		20	5.9	
Self Serve	5	1	0.5			0.5	0.25	1	0.5	0.25	0.5		1	1		31.5	9.3	
AI	4	1	1			1	1	1	1	1	1	0.25	1	1	x	22	6.5	
Intimacy	4			1		1	1	1	1	1	1		1	1		14.5	4.3	
Escalation	3					1	0.25	0.25	0.5	0.5	1		1	1	1	27.25	8.1	
EX	3	1	1	1		0.5			1	1		0.5	0.5	1	0.5	22	6.5	
IOT	3	1	1			1	0.5		1	1			1	1		10	3.0	
Loyalty	3					1	0.5			1			1	1		32.5	9.6	
Prediction	3	1	1	1		1	1	1	1	0.5	1	1	1	1	0.5	33.75	10.0	
Automation	3	1	1		0.5	1	1	1	1	0.5	1	0.25	0.25	1	0.5	5	2.1	
Mobile	2		0.25		0.25											21	8.8	
Proactivity	2	1	0.5	1		1	0.5	0.5		1			1	1	0.5	8	3.3	
Causes	2		0.25							0.25			1			13.5	5.6	
Videos	2	0.5					1	0.5			1	0.5	1	1		19.5	8.1	
Voice	2	0.5	0.25			1	1	1			1	1	1	1		7	2.9	
Communities	1		0.25	0.5			0.25	0.25	0.25	0.25		0.25	0.25			7.25	3.0	
Free offers	1			0.5			0.25	0.25		0.5			1	1		20	8.3	
SM messages	1	1				0.5	1	1	1	0.5	1		1	1		24	10.0	
SMS	1	1				1	1	1	1	0.5			1	1	1			

The horizontal axis of the graph was built by combining the importance and relevance of AI trends for each CX trend. AI trend importance was derived in a

similar way to CX trend importance: The first 6 websites returned from the AI search with useful information were analysed for trend indications. This revealed 14 most significant AI trends. The importance of each trend was the count of the number of sites mentioning that trend. The relevance of an AI trend to each CX trend was based on a subjective assessment of how much it supported the CX trend in question. This ranged from 0 for no support, through 0.25 or 0.5 for partial support to 1 for full support. For each CX trend the products of AI trend importance and trend relevance were summed. This gave a raw score for overall AI support of each trend. This raw score was higher for CX trends with more supporting AI trends, with more important AI trends and with more relevant AI trends. The raw AI support scores were normalised to get a scale between 1-10 for the graph.

The graph above outlines three distinct cohorts of CX initiatives and a frontier line. The first cohort at the top of the graph shows Omnichannel CX and Chatbots as being the most important CX trends mentioned in high-ranking online searches. The second cohort on the right of the page identifies Personalisation, general AI usage, Automation and Predictive Analysis as relatively important CX trends that are most supported by the ongoing evolution in AI now. The bottom left cohort including Escalation to Human Agents, Video Marketing, Loyalty Programs, Causes Related CX, Mobile UX, Communities of Engagement and Free Offers in Online Marketing are all probably more suited to specific niche plays because they lack the momentum of other CX trends. The other CX trends outside of these cohorts are harder to cluster. They do indicate a frontier of trends likely to progress at similar rates. This idea of temporal evolution is expanded in the graph below.

FIGURE 12: WAVES OF AI SUPPORTED CX



The figure above shows the same CX trends clustered temporally into five waves of evolution. This suggests Chatbots and Personalisation will be in the first wave. Omnichannel, AI based CX, Automated CX and Predictive Analytics are expected to comprise a second wave and so on.

- ChatGPT and Varian’s Rule

This book was written using methods likely to become more commonplace in the future: the author collaborated with ChatGPT over approximately 4 weeks. The productivity uplift from this collaboration was remarkable. The last two books written by the author took 12-18 months to complete. All three books were written as side hustles to the author’s full time consulting work. The form of the collaboration with ChatGPT was simple. The author worked out what the book was supposed to cover and then prompted ChatGPT to create most of the content. Sometimes the queries to the LLM were iterated to get good enough detail. Other times they were completed directed by the author’s existing knowledge or secondary research. The resulting text was then edited at three levels – by the MS Word spellchecker, with a Grammarly plug-in and then with humans manually reviewing and editing the text. ChatGPT did not create content from its trained knowledge base. It only knows essentially a summary of the Internet up to 2021. Instead the WebchatGPT plugin for Chrome was installed. This plug in directs ChatGPT to search for current web pages to any specific content question and asks the LLM to summarise the search results with references. In other words, the LLM was only wordsmithing not actually generating novel semantic meaning. This approach was needed to avoid the problems of so-called hallucinations.

Hallucinations are when the response an LLM gives is false because it is filling in the gaps in its knowledge to construct what looks like a sensible block of text. LLMs don't really know anything as concrete facts. What they know is what words are likely to go together in a way a human would recognise as a complete and understandable communication. They are literally trained on what word to offer next given a stream of words. By doing this over and over they can generate sophisticated responses approaching 'real and factual' responses. But they are not trying to speak the truth because they have not been trained to do so. An LLM's confident inaccurate response is called a hallucination.

LLMs hallucinate because during training they summarise large amounts of text into something like connected chunks of meaning. This is similar to how humans seem to code episodic memory. When we remember an event, say like the last time we saw a family member or what happened during a job interview, we don't store the event as a complete movie with all our senses. Instead our Reticular Activating System and our Hippocampus seem to identify significant parts of the event and only store them. We might not remember some obvious details like the colour top the other person wearing or the way they were wearing their hair. We do remember seeing them because this was significant. We may not remember word for word exactly what was said. We probably do remember the meaning we took away. We might even be able to quote key sound bites. It almost as if we record only significant bullet points of events as memories. Later when we try to recall what happened, we reconstruct the episode as a kind of movie in our mind from these summary significances. The critical aspect of this type of memory is we don't seem to record everything. Instead we only capture what seems to be important at the time. Episodic memory recall in humans can present several accuracy challenges, which can serve as a metaphor for explaining why large language models can sometimes hallucinate. This type of memory tends to be highly subjective and influenced by individual perspectives and experiences. Each person's recollection of an event is unique to them¹⁹⁹. Similarly, large language models generate responses based on the vast amount of data they have been trained on, which can include subjective information. This subjectivity can sometimes lead to hallucinations or false information being generated. Human episodic memory is also susceptible to distortion over time. Memories can be altered or modified during recall, leading to inaccuracies or false details being incorporated into the recollection. Factors such as personal biases, external suggestions, and the passage of time can contribute to memory distortion²⁰⁰. Similarly, large language models can generate information based on patterns and associations they have learned from training data, which can include distorted or incorrect information. These distortions can result in hallucinations or the generation of inaccurate content. Human sometimes find their memories

cannot be fully retrieved or recalled. This can occur due to various factors such as interference from other memories, decay over time, or insufficient retrieval cues. Incomplete recall may result in gaps or missing details in the recollection of an event²⁰¹. Large language models might not have access to the complete information needed to generate a response accurately. In such cases, they may attempt to fill in the gaps based on incomplete or partial information, leading to the generation of speculative content. Finally human episodic memory can be influenced by cognitive biases, including confirmation bias, where individuals tend to remember information that confirms their existing beliefs or expectations. This bias can affect the accuracy and objectivity of memory recall²⁰². Large language models can also exhibit biases based on the data they have been trained on, leading to biased or skewed information being generated. These biases can contribute to the occurrence of hallucinations or the propagation of misinformation.

Prompting ChatGPT to build longer text from the author's significant bullet points was one way to avoid hallucinations. The other way was to supply web pages to the LLM and have it summarise what was on those pages. This was where the WebChatGPT plug in was a very useful tool. This plug in takes a query and adds search results to prime the LLM to create a response. Both techniques for using ChatGPT as a collaborator do not rely on its less granular understanding of topics from its pre 2021 internet training.

Using an LLM as a writing collaborator is an example of Varian's Rule in action. A wealthy author could have a personal assistant who responded to their verbal dictation and completed search requests to craft text to save the actual author time. ChatGPT with the WebchatGPT plugin operated much the same way. However this software based approach was essentially cost free and much faster than having a personal assistant. Arguably a smart, motivated personal assistant could offer higher quality research and text, but ChatGPT was quite acceptable for the most part. None of the conceptual content was developed by the LLM except as explicitly stated in the copy. It could not produce the trend analysis work shown above in this chapter for example. However the specific words of the whole book were probably 80% crafted by the LLM. Here is an example of a typical query interaction originally intended to create the content for this section of this chapter.

User: describe the Varian Rule and use it to predict the trends in customer experience (CX) most likely to occur in with the greatest impact in the future.

ChatGPT4.0: I'm sorry, but I can't find any information on the "Varian Rule" in the provided search results. If there is an established principle or theory called the Varian Rule, it's not mentioned or explained in the references you've provided. If you

can provide more details or specific sources mentioning the Varian Rule, I would be more than happy to help elaborate on it and apply it to predicting trends in customer experience (CX).

Nonetheless, I [response interrupted by user]

At this point the author decided it needed to be presented with Varian's Rule and then maybe it could create a better answer. Over time the author learned to treat the LLM as if it is a motivated and smart intern who is neuro atypical. Sometimes it just doesn't quite understand nuance. In that case the problem is with the author not the intern LLM. In software terms this has been called a PEBCAK error (Problem Exists Between Chair And Keyboard) or more simply coded as ID10T. Here is the next go at the query which explains Varian's Rule and asks for some implications related to CX trends. Having to go into this level of detail was strange because the text on Varian's Rule came from a simple web search that turned up the original Financial Times article by Andrew McAfee²⁰³. Normally the WebchatGPT would have been expected to find this (or the relevant Wikipedia article) and use it without needing such additional detail.

User: Google's chief economist Hal Varian, proposed an elegant and thrifty approach to prediction. "A simple way to forecast the future," he says, "is to look at what rich people have today." This works. Applying this method a few years ago would have led one to foresee the rise of Uber and the spread of smartphones around the world, to take just two examples. Hal's point is that tech progress quite quickly makes initially expensive things — both goods and services — cheaper, and so hastens their spread. Which is why this progress is the best economic news on the planet. This idea is called Varian's Rule. Please try to apply Varian's rule to predict which of the most important trends in customer experience (CX) are more likely to manifest faster.

This time the WebchatGPT plugin had more success in finding some search results. It reworked the query above to become the whole of the quoted text that follows. You can see it started by finding 8 search results (it was allowed to use up to 10) and then it crafted a very specific query. You can see this in the instructions section which included my original query above to send off to ChatGPT4.0.

Web search results:

[1] "Apr 7, 2015 · Applying this method a few years ago would have led one to foresee the rise of Uber and the spread of

smartphones around the world, to take just two examples. Hal's point is that tech progress ..."

URL: <https://www.ft.com/content/4329a987-9256-3059-b36f-1aba9338b800>

[2] "Dec 20, 2017 · The idea is that you can forecast the future by looking at what rich people have today, and then in another decade or so the middle class will have it, and in two decades the poor will have it."

URL: <https://www.aei.org/economics/googlenomics-a-long-read-qa-with-chief-economist-hal-varian/>

[3] "Apr 7, 2015 · Applying the method a few years ago would have led one to foresee the rise of Uber and the spread of smartphones around the world, to take just two examples. Mr Varian's point is that tech ..."

URL: <https://www.ft.com/content/5a484a28-dd2c-11e4-a772-00144feab7de>

[4] "Aug 16, 2017 · Hal Varian '69. Google's chief economist uses math to model behaviour. By. Elizabeth Karagianis. August 16, 2017. Hal Varian was raised on a farm in a sleepy Ohio town. Growing up, he turned to ..."

URL:
<https://www.technologyreview.com/2017/08/16/242148/hal-varian-69/>

[5] "Free tools like Google Correlate and Google Insights for Search enable you to use that database of intentions to 'predict the present' and better understand your customers' behaviour in real time. I recently asked a group of Googlers which day of the week had the most Google searches for the word 'hangover'. Most of them chose Sunday or ..."

URL: <https://www.thinkwithgoogle.com/marketing-strategies/search/predicting-the-present/>

[6] "Jun 5, 2018 · Varian is certainly right that the dearth of data analysts is acute—IBM refers to it as a "quant crunch. "According to data from job website Indeed, data scientist job postings in the US ..."

URL: <https://qz.com/1296930/hal-varian-googles-chief-economist-thinks-the-world-needs-more-data-scientists>

[7] "Dec 15, 2017 · Google's chief economist says that, although the advent of machine-learning techniques using big data has provided important new insights into how the real economy reacts to certain economic stimuli, it is likely there are multiple factors at work, rather than any single economic principle. In an interview with Central Banking, published on ..."

URL: <https://www.centralbanking.com/central-banks/economics/3335411/google-chief-economist-hal-varian-offers-machine-learning-insights>

[8] "A discussion with Google's Hal Varian. Google's Chief Economist, Hal Varian, talks to Thomson Reuters Global Business Director, Asif Alam, about the secrets to Google's success, their biggest bets, big data and partnership, and where the global economy is headed. 08:01. Discover Thomson Reuters technology solutions."

URL: <https://www.thomsonreuters.com/en/products-services/technology/partner-stories-googles-hal-varian.html>

Current date: 17/05/2023

Instructions: Using the provided web search results, write a comprehensive reply to the given query. Make sure to cite results using [[number](URL)] notation after the reference. If the provided search results refer to multiple subjects with the same name, write separate answers for each subject.

Query: Google's chief economist Hal Varian, proposed an elegant and thrifty approach to prediction. "A simple way to forecast the future," he says, "is to look at what rich people have today." This works. Applying this method a few years ago would have led one to foresee the rise of Uber and the spread of smartphones around the world, to take just two examples. Hal's point is that tech progress quite quickly makes initially expensive things — both goods and services — cheaper, and so hastens their spread. Which is why this progress is the best economic news on the planet. This idea is called Varian's Rule. Please try to apply Varian's rule to predict which of the most important trends in customer experience (CX) are more likely to manifest faster.

The result back from ChatGPT4.0 is close to miraculous because when you review the search results they only relate to Varian's Rule. The rest of the content in the response below comes from what ChatGPT summarised from its review of the entire internet up to 2021. This means it already understood enough about where CX was going in 2021 to produce the response below. The numbers in square brackets [] relate to the references in the modified prompt created by the WebchatGPT plugin above. They are very useful for fact checking what the LLM is asserting.

ChatGPT4.0: Varian's Rule, as proposed by Google's chief economist Hal Varian, suggests that to predict the future, one can look at what the wealthy have today [2]. This rule is based on the observation that technological progress often makes initially expensive goods and services cheaper, thereby hastening their adoption among a wider segment of the population.

Applying Varian's Rule to customer experience (CX) trends, we could make the following predictions:

Personalization: Wealthier individuals are already enjoying highly personalized experiences, with services and products tailored to their needs and preferences. As technology continues to evolve, such personalization will become more affordable and accessible to all, thereby becoming a widespread trend in CX [8].

Predictive Experiences: Companies are increasingly using big data and predictive analytics to anticipate customer needs and provide proactive service. This trend is currently more prevalent among high-end businesses that can afford these technologies, but as per Varian's Rule, it will likely trickle down to smaller businesses and become a standard in CX in the future [1, 2, 5].

AI-Driven Customer Service: Wealthy customers often have access to round-the-clock personal customer service. With the advancement of AI and machine learning technologies, this level of service will become increasingly cost-effective, enabling businesses to provide 24/7 customer service to all their clients. This will likely lead to a surge in AI-driven customer service experiences [6].

Direct-to-Consumer (DTC) Services: The trend of direct-to-consumer services is increasing, with customers preferring to buy directly from manufacturers or service providers. This has

been more common in high-end sectors but is likely to become commonplace across all industries, enhancing the overall customer experience [6].

Remember, Varian's Rule is a general guideline and not a hard and fast rule, and technological advancements and societal factors will play a significant role in shaping these trends.

If you compare the LLM response (which took about 3 minutes of the author's time to create) with the start of this chapter (which consumed more like 8 hours with all the research and graphs), you can see the software did a pretty good job the second time around. It has missed Omnichannel and Chatbots, probably because they were not as relevant in 2021. It has added in Direct to Consumer. This is probably still a hallucination because reference [6] has nothing in it about direct-to-consumer commerce. The article relates to quotes from Varian about a lack of data scientists. This example shows both the capabilities and limitations of LLMs. It also points out the value of getting the LLM to provide references. The quality of the output is highly dependent on the quality of its directions, including the base content it is asked to consider. Much like any smart, motivated intern.

• Potential Scenarios

It is useful to propose some potential future scenarios to frame context of how CX will develop with AI support. The seven scenarios presented start with incremental developments possible simply with more time and money being invested. These relate to computing power and healthcare applications. Evolutionary developments likely to follow are based on chatbots becoming smarter, more widespread, and more human. The most far off transformational developments close the human-machine gap from the other side. The prospect of cyborgs seems like science fiction even though in many limited ways many humans are already cybernetic. These scenarios are listed in order of likelihood, not impact. The range of scenarios is not necessarily comprehensive because there is limited data available to build them on.

More Powerful Hardware on Offer

Artificial Intelligence models, especially large ones, can require substantial computational resources, often provided by high-performance systems or supercomputers, particularly during their training phase. For instance, OpenAI has used a supercomputer with more than 285,000 CPU cores, 10,000 GPUs, and significant network connectivity for each GPU server for the development of its AI models²⁰⁴. This supercomputer ranks among the top five in the world, according to Microsoft. The cost and resources required to train large AI models, such as generative AI, are significant. For example, training a large

language model like ChatGPT, which has 175 billion parameters, necessitates considerable resources²⁰⁵.

Once trained AI models typically require less computational power for the inference phase (when the model is applied to new data to generate predictions or responses). Supercomputers or high-performance computing systems may be necessary for training sophisticated AI models initially. But then running the trained models does not necessarily require a supercomputer. Once trained an LLM can be operated on much more modest hardware.

Cloud computing platforms such as Microsoft's Azure, Amazon Web Services (AWS), and Google's Cloud Platform are not considered supercomputers. Instead, they use vast networks of regular servers distributed across multiple locations around the globe, each of which is a data centre. Cloud computing is a model where servers, networks, storage, development tools, and even applications are enabled through the internet. This offers faster innovation, flexible resources, and economies of scale.

While these platforms are not supercomputers in themselves, they can be used to build virtual supercomputers by harnessing the combined power of their distributed servers to perform tasks that traditionally required supercomputing power. In this way, the lines between cloud computing and supercomputing are somewhat blurred.

The market for cloud computing has been experiencing significant growth and is expected to continue to do so²⁰⁶. As of 2022, the global cloud computing market was valued at USD 569.31 billion and is projected to grow to USD 2,432.87 billion by 2030, exhibiting a CAGR of 20.0% during the forecast period. Cloud services are becoming increasingly diversified, with organizations increasingly opting for multi-cloud environments that take advantage of different cloud models, both on-premise and off-premise.

In terms of the supply-demand dynamics, there's no clear indication of the market being either undersupplied or oversupplied. However, given the high growth rates and the ongoing digital transformation of industries globally, the demand for cloud services is strong and likely to remain so for the foreseeable future. The market is predicted to grow at a robust CAGR of 13.7%, reaching nearly US\$ 1.9 Trillion by 2032²⁰⁷. This suggests that the supply of cloud services will need to continue growing to meet this increasing demand. The ease of access and lower costs associated with cloud computing are enabling both mature and new companies to leverage these services. This democratization of access to powerful computing resources is likely to further drive demand in the future.

Graphcore is a startup based in the UK that focuses on developing AI chips for next-generation computing²⁰⁸. The company aims to shape the future of AI and business strategy by providing powerful hardware solutions for artificial intelligence applications. The firm was founded in 2016 by Nigel Toon and Simon Knowles. The company's goal is to create a new generation of silicon chips specifically designed to power AI programs, enabling more efficient and powerful machine learning and deep learning algorithms. Graphcore envisions a future where their AI technology contributes to a democratized intelligence that benefits everyone.

Graphcore has attracted significant attention and investment. The company has raised substantial funds, including a 2020 investment that brought its market valuation to over \$2 billion²⁰⁹. Graphcore's main focus is on developing its proprietary Intelligence Processing Unit (IPU). These chips are specifically designed to accelerate AI workloads, providing higher performance and energy efficiency for machine learning and deep learning tasks. The IPUs aim to address the growing demands of AI applications, including voice recognition, self-driving cars, and more. Graphcore's AI chips put the company in direct competition with industry giants such as Intel and Nvidia. By offering specialized hardware solutions for AI, Graphcore aims to provide alternatives to existing technologies and advance the field of artificial intelligence. In addition to developing AI chips, Graphcore offers cloud services that enable enterprises to harness the power of IPUs in their applications. These services include pre-trained models, optimized inference engines, and APIs to build and deploy AI-native products and platforms.

Graphcore has announced plans to develop an AI supercomputer known as the Good Computer. The company claims that this supercomputer will be one of the most powerful of its kind, further enhancing their position in the AI hardware market. They faced a setback in 2020 when Microsoft reportedly scrapped an early deal to purchase Graphcore's chips for its Azure cloud computing platform.

Quantum Computing is a Flop

The current state of quantum computing can be described as a mixture of progress, breakthroughs, and ongoing development. While there have been significant advancements in the field, there are also criticisms regarding the hype surrounding quantum computing. Quantum computing aims to harness the principles of quantum theory, which govern the behaviour of energy and matter at atomic and subatomic levels. By utilizing qubits, quantum computers can store and manipulate information in ways that classical computers cannot.

A qubit, short for "quantum bit," is the fundamental unit of information in quantum computing, analogous to the classical bit in traditional computing. It

is the basic building block of quantum information and is implemented using a two-state quantum system that follows the principles of quantum mechanics. Unlike classical bits, which can only exist in one state (either 0 or 1), qubits can exist in a superposition of both states simultaneously, allowing for more complex and powerful computations.

In quantum mechanics, a qubit can be represented by various physical systems, such as the spin of an electron or the polarization of a photon. The two states of a qubit can be interpreted as spin up and spin down for an electron or vertical and horizontal polarization for a photon. However, it's important to note that qubits are not limited to these specific physical realizations and can be implemented in different ways depending on the technology used. The unique characteristic of a qubit is its ability to be in a coherent superposition of both states, thanks to the principles of quantum mechanics. Ideally this superposition enables quantum computers to perform parallel computations and solve certain problems more efficiently than classical computers.

The practical implementation of qubits varies across different quantum computing platforms, including superconducting qubits, trapped-ion qubits, topological qubits, and more. Researchers and engineers are continuously working on developing new technologies to improve qubit coherence, reduce errors, and scale up the number of qubits to build more powerful quantum computers.

Qubits unique properties make quantum computers powerful tools for solving complex problems in various fields. The field has witnessed notable achievements over the past few years. According to a survey conducted by McKinsey & Company, technology, media, and telecom companies have made several breakthroughs, including achieving quantum supremacy, developing industrial quantum computers, and establishing cloud-based solutions²¹⁰.

Quantum computing also faces several significant challenges and problems that need to be addressed for its widespread adoption and practical application. Here are some of the major issues currently associated with quantum computing²¹¹:

Scalability: One of the biggest challenges in quantum computing is achieving and maintaining scalability. While small-scale quantum systems with a few qubits have been developed, scaling up to larger numbers of qubits while maintaining their coherence and reliability remains a significant hurdle. It is essential to overcome this challenge to demonstrate the potential power of quantum computing for solving complex problems.

Decoherence and Quantum State Preservation: Quantum systems are susceptible to decoherence, which refers to the loss of quantum information

due to interactions with the environment. Maintaining the coherence of qubits over a sufficiently long period of time, known as quantum state preservation, is crucial for performing accurate computations. Extending the coherence times of qubits is a major challenge in quantum computing.

Quantum Error Correction: Quantum systems are highly sensitive to noise and errors caused by environmental factors and interactions with their surroundings. Quantum error correction techniques are crucial to mitigate these errors and protect the delicate quantum information stored in qubits. Developing effective error correction methods that can address both physical and logical errors is an ongoing research area.

Qubit Connectivity and Gate Operations: Qubits need to be connected and interact with each other to perform computations. Ensuring high-fidelity and reliable gate operations between qubits is critical for implementing quantum algorithms. Achieving long-range qubit connectivity and improving the accuracy of gate operations are ongoing research goals.

Quantum Supremacy and Application Development: Quantum supremacy refers to the demonstration of a quantum computer solving a problem that is infeasible for classical computers. While quantum supremacy has been achieved for specific tasks, such as simulating quantum systems, general-purpose quantum computers capable of outperforming classical computers on a wide range of practical problems are still being developed. Bridging the gap between quantum algorithms and real-world applications is a significant challenge.

Cost and Hardware Constraints: Quantum computing hardware is currently expensive to develop and maintain. The creation of stable and reliable quantum systems requires specialized and highly controlled environments. Overcoming cost and hardware constraints is crucial for making quantum computing more accessible and commercially viable.

All this suggests quantum computing is not likely to disrupt the current cloud platforms or super computers any time soon.

Better Healthcare Saves 1% of GDP

There have been several studies estimating the potential cost savings that could be achieved through the application of AI in healthcare. According to research from Harvard University and consulting firm McKinsey, the use of AI technology could save the US healthcare industry up to \$360 billion per year, which represents about 10% of annual healthcare spending²¹².

It is important to note, however, that these estimates are based on current technology and are subject to change as AI technology advances and as the healthcare industry adapts to incorporate AI in more areas. AI can support

improvements in care outcomes and patient experience, transforming how healthcare is delivered. For instance, AI could enable truly proactive, predictive healthcare, allowing health systems to intervene even before the disease occurs or progresses.

Freenome is an innovative startup based in the United States that operates at the intersection of biology, machine learning, and medicine²¹³. Their primary focus is on developing diagnostic tests for the early detection of cancer and other diseases using a liquid biopsy approach, which involves analyzing blood samples. By leveraging artificial intelligence (AI) algorithms, Freenome aims to assist doctors in detecting cancer at its early stages and enabling more effective treatment strategies.

What sets Freenome apart is their commitment to transparency regarding the limitations and potential biases of their algorithms. They recognize the importance of addressing these concerns and ensuring the reliability of their testing methods. By combining machine learning techniques with biological data, Freenome aims to spot patterns that can help identify cancer and make informed treatment decisions. With their mission to empower individuals to proactively address cancer and other diseases, Freenome is developing non-invasive and accessible disease screening products that provide accurate results for early-stage cancer detection.

AI can increase the productivity and efficiency of care delivery, which will allow healthcare systems to provide better care to more individuals. For instance, AI can free up millions of hours annually by leveraging applications such as imaging, wearables, virtual health assistance, and robotics.

Viz.ai is a pioneering startup based in the United States that harnesses the power of artificial intelligence (AI) to enhance stroke diagnosis and treatment while fostering trust between patients and healthcare providers²¹⁴. Their innovative approach revolves around using cutting-edge technology to streamline clinical workflows and coordinate care more effectively, ultimately reducing systemic delays that could impede life-saving treatments.

At the core of Viz.ai's offerings is their suite of AI-powered products designed to detect and alert care teams to suspected diseases within their network. Their advanced algorithms and machine learning techniques have been instrumental in increasing the speed of diagnosis and care, covering a substantial population of over 220 million lives across more than 1,300 hospitals and health systems. By utilizing their platform, care teams can connect earlier, facilitating expedited diagnosis and appropriate care. This optimized clinical care pathway ensures that patients receive timely access to the right specialist, significantly improving the overall quality of care delivered.

AI applications in healthcare, such as patient self-service, chatbots, and computer-aided detection systems for diagnosis, can reduce costs and errors while increasing convenience and efficiency. AI plays a critical role in fostering advancements in the healthcare industry, including the productivity of pharmaceutical companies. AI can help streamline the process of creating, testing, and receiving approval for new drugs, potentially saving critical time for patients.

Chatbots Get Even Smarter

Chatbots are likely to get smarter in the future by becoming more integrated into marketing and customer service strategies, as well as benefiting from improvements in mobile technology and AI advancements. For instance, improvements in smartphone processing power, along with an increasing number of chatbot apps, will likely enhance the capabilities of AI chatbots, especially on mobile platforms. AI-driven chatbots, are revamping the ways of two-way communication between businesses and customers. This includes machine learning capabilities that allow for more flexible automated business communications.

Hugging Face, based in Paris, France, is an open-source and platform provider of machine learning technologies²¹⁵. The company was founded in 2016 and has gained recognition for its contributions to the field of natural language processing (NLP). Hugging Face's notable contribution is the development of the transformers library, which has become widely adopted in the NLP community. The library provides a range of pre-trained models that can be fine-tuned for various NLP tasks. Furthermore, Hugging Face's platform serves as a hub where users can share machine learning models and datasets, facilitating collaboration and knowledge exchange within the machine learning community.

One of Hugging Face's significant achievements is the creation of BLOOM, a language model with an impressive 176 billion parameters. BLOOM has the capability to generate text in 46 natural languages and 13 programming languages, making it the largest language model in terms of parameters for many of these languages. The development of BLOOM involved the collaboration of over 1,000 researchers from 70 organizations, culminating in this groundbreaking advancement in language modelling.

It's difficult to make a direct comparison between Hugging Face's Bloom and OpenAI's ChatGPT-4.0. Both models are advanced large language models that have been trained on extensive data. OpenAI's GPT-4.0 is a multimodal model, which means it can accept both image and text inputs, and it exhibits human-level performance on various professional and academic benchmarks²¹⁶. It is also capable of generating, editing, and iterating on creative and technical writing tasks, such as composing songs or writing screenplays. Bloom by

Hugging Face is an autoregressive LLM that's been trained on a massive amount of text data. It's capable of outputting coherent text in 46 languages and 13 programming languages, making it highly versatile and adaptable.

GPT-4.0's larger size and multimodal abilities could make it more versatile for certain tasks, while Bloom's multilingual capabilities could make it more suitable for tasks involving multiple languages.

Given that chatbots are predicted to handle most customer service interactions within the next 10 years, it's plausible that chatbot technology will become sophisticated enough to engage in human-like conversations, thereby allowing for in-depth assessment of human responses. Such a development would likely involve improvements in the chatbot's ability to process human language, create highly personalized experiences, and engage customers effectively.

However, this development could also present certain challenges. For instance, human-like chatbots, while they can provide contextually appropriate responses, may lack the originality and spontaneity of human responses.

In this possible future, companies like Receptiviti could use these advanced chatbots to assess and analyse human communication for various purposes. This could range from gauging customer sentiment and satisfaction to even providing psychological insights, based on the patterns and nuances of a person's interactions with the chatbot.

Receptiviti (Canada) is a startup that uses natural language processing and psychological profiling to personalize customer interactions²¹⁷. Their computational language psychology platform helps businesses understand the emotions, drives, and traits that affect human behaviour. Receptiviti's proprietary algorithms uncover signals from everyday human language that traditional NLP techniques do not, providing explainable, out-of-the-box insights with minimal training required.

Receptiviti makes money by offering their platform to businesses, which use the insights gained from analyzing language data to predict customer behaviour, improve customer experiences, and inform decisions. By focusing on the psychological aspects of language, Receptiviti can provide unique insights into an individual's emotions, drives, and traits, which can help businesses better understand their customers and tailor their interactions accordingly. Customers use Receptiviti to analyse various types of data, such as social media, news, job applications, and essays, to uncover insights about people that can power predictive models and optimize interpersonal interactions. By using AI's natural language processing capabilities, Receptiviti

can automatically mine and analyse customer data to improve customer satisfaction and experiences.

Chatbots are Everywhere

In a future scenario where chatbots are ubiquitously available, we can envisage their presence in multiple facets of our daily lives, from personal use to business applications, and even in areas such as healthcare, education, and public services.

On a personal level, chatbots could become our everyday virtual assistants, present in our phones, computers, smart speakers, and even our cars. They could help us schedule appointments, keep track of our to-do lists, provide weather updates, and answer our random queries, much like Siri, Google Assistant, and Amazon Alexa do today but with a higher degree of sophistication and personalization.

In the business sector, chatbots could be widely deployed for both customer-facing and employee-facing scenarios. For customers, chatbots could handle most inquiries, provide updates on orders, process returns or exchanges, and even suggest products based on customer preferences. For employees, chatbots could help answer common IT queries, provide technical support, and assist in various other tasks, thereby improving productivity. In the field of marketing, chatbots could play a significant role in content marketing by collecting and organizing valuable information from people interested in a brand, which could then be used for more precise target marketing.

In the retail industry, chatbots could assist with customer inquiries, potentially saving billions of dollars and a significant amount of time on customer service inquiries. In the healthcare sector, chatbots could help manage appointments, provide medication reminders, and answer common health-related queries. For education, chatbots could assist students with learning materials, answer questions, and provide personalized study plans.

MindMeld is a conversational AI platform developed by a startup acquired by Cisco for \$125 million²¹⁸. It is a Python-based machine learning framework designed to create production-quality conversational applications. The platform provides tools and capabilities for every step in the workflow of a state-of-the-art conversational application, including advanced natural language processing.

The MindMeld platform has evolved over several years and is used in various applications, such as chatbots and interactive voice response (IVR) systems]. It consists of a natural language processor that analyses user inputs through a hierarchy of classification models, each helping the next tier of models by narrowing the problem scope. Apart from the machine learning framework, Cisco has also released extensive documentation and a step-by-step guide

called the Conversational AI Playbook²¹⁹. This playbook serves as a practical guide for applied machine learning engineers to build modern, state-of-the-art conversational experiences. It includes the following sections covering 'Building a Conversational Interface in 10 Steps':

Step 1: Select the Right Use Case

Step 2: Script Your Ideal Dialogue Interactions

Step 3: Define the Domain, Intent, Entity, and Role Hierarchy

Step 4: Define the Dialogue State Handlers

Step 5: Create the Knowledge Base

Step 6: Generate Representative Training Data

Step 7: Train the Natural Language Processing Classifiers

Step 8: Configure the Language Parser

Step 9: Optimize Question Answering Performance

Step 10: Deploy Trained Models

The playbook provides a comprehensive understanding of MindMeld's capabilities and guides users through the process of creating advanced conversational AI applications. It covers essential concepts, best practices, and methodologies for building effective and efficient conversational experiences. This means anyone who wants to go through the learning process can design, build, and deploy a chatbot right now.

User Interfaces Become More Human

Soul Machines, based in Auckland, New Zealand, is a pioneering deep technology firm that has made significant strides in the field of generative AI and autonomous animation²²⁰. Their work primarily focuses on the creation of extremely lifelike digital avatars or "digital humans," which are designed to provide compelling, personalized experiences for online audiences. These avatars are designed to emotionally react to users, offering a novel level of interaction that makes the experience feel more human-like.

Their technology is so advanced that it can create a digital person in real-time, allowing for more personalized and empathetic interactions. The company has been successful in creating numerous humanoids that have a striking resemblance to real people, complete with wrinkles and moving eyes. This innovation has been utilized by several well-known brands, celebrities, and influencers, to provide a unique customer service experience. Soul Machines' AI-powered avatars have been well-received by users, with over 81% of

customers expressing interest in conversing with these avatars again and 89% claiming that they achieved their goals.

In a future scenario where user interface becomes more humanised, one can envision a symbiotic relationship between humans and technology, where the lines between digital and physical reality become increasingly blurred. This could be facilitated by advanced AI systems like those currently being developed by companies like Soul Machines.

Imagine, for instance, a digital personal assistant powered by AI, designed to emulate human behaviour, look, and sound like a real person, and understand and respond to emotions. This assistant, created by Soul Machines, could interact with users through a screen or VR headset, providing a realistic, emotionally engaging, and empathetic experience. The assistant could use voice technology like Speechelo, making interactions feel more natural and personalised²²¹. Users could carry out tasks like setting appointments, searching the internet, or getting personalized recommendations, all through engaging in lifelike conversations with the AI. This could create a more human-like interaction, potentially making technology more accessible and intuitive, especially for those who may currently struggle with more traditional UI paradigms.

Such a future could change the way we interact with technology on a daily basis, promoting a more human-centric design approach and redefining the concept of human-computer interaction. It could also have significant implications for sectors like customer service, healthcare, and education, among others, by providing users with more personalised, empathetic, and engaging experiences.

The next wave of digital tech, often referred to as "smart tech," has the potential to help us rehumanize work and daily activities, and voice-guided user interfaces will play a significant role in this development. Voice-guided user interfaces are predicted to have an adoption rate of over 80% in the next five years. This technology will be crucial as our screens become smaller, making interfaces harder to navigate. While touch will likely still dominate mobile devices for a while, voice interaction promises a deeper and more ubiquitous method for interacting with technology. This is where Speechelo-like systems come in, enabling a highly natural and intuitive voice interaction that goes way beyond setting appointments and surfing the web.

Speechelo is an AI text-to-speech software that allows users to convert written text into natural-sounding speech. It offers a range of features and benefits that make it a useful tool for various applications such as video narration, podcasting, and e-learning content creation. The software provides over 30 human-sounding voices in multiple languages, allowing users to choose from different accents and tones to find the perfect voice for their

projects. One of the notable advantages of Speechelo is its ability to add inflections to the generated voice, making it sound more expressive and natural. Speechelo claims to deliver high-quality voiceovers that sound 100% human, setting it apart from many other text-to-speech solutions.

In this future scenario, imagine using your voice to command your smart devices at home, at work, or even in public spaces. You could ask your fridge what ingredients you have and what you can cook with them or request your car to plan a route with the least traffic. You could have a conversation with your personal AI assistant, who could help you manage your schedule, answer your queries, or even engage in meaningful conversations to help combat loneliness.

Simultaneously, other technologies like Virtual/Augmented Reality (VR/AR) and wearables will also contribute to a more humanised UI. VR/AR could render TVs obsolete, providing immersive experiences directly in your living room, while wearables could offer unique interactions that go beyond the confines of screens.

Humans Become Cyborgs

In a potential future scenario, we could become cyborgs due to the development of implantable brain-computer interfaces (BCIs) by companies such as Neuralink. Neuralink is a neurotechnology company that focuses on developing implantable brain-computer interfaces (BCIs)²²². The primary aim of the company is to develop technology that could eventually help disabled patients move, achieve "symbiosis" between the human brain and artificial intelligence, and enhance human cognitive abilities. These BCIs could transform our interaction with technology and the world at large in unprecedented ways.


A glimpse into this future can be seen from Neuralink's previous work. In 2020, the company implanted a computer chip into a pig's brain, sparking conversations about the possibilities and dangers of brain-computer interfacing²²³. Further refinement of this technology could allow it to be safely applied to humans, leading to a new era where our brains can directly interface with digital devices. This would allow our brains to directly control devices around us. For instance, we could potentially turn our car on with our mind, control the thermostat in our homes, send and read text messages, or even play video games without the need for physical controllers.

This would be an extension of the current trend where millions of people already live with technology in or on their bodies. From dental implants to pacemakers and from hearing aids to smart glasses, humans are increasingly integrating technology into their bodies. However, becoming cyborgs doesn't necessarily mean instant extraordinary abilities. The integration of technology

into our bodies could be a gradual process, and the capabilities that we acquire would be shaped by the limitations and possibilities of the technology itself, as well as our capacity to adapt to and utilize it.

This could lead to a symbiotic relationship between humans and AI, driving a win-win situation that combines human creativity and emotional intelligence with the analytical power and speed of artificial intelligence. Neuralink is still in the research and development phase, with clinical trials in humans being prepared from early 2022. Its current applications are primarily centred around medical advancements and enhancing human capabilities.

X. Conclusion:



Fortunes changed after AI took over everything.

This book has been created to change the paradigm about automation and CX. Historically firms have not chosen to automate their CX because it was better for their customers (or if they asserted they did they were likely in error). Historically the only advantage of automated systems for customers is immediate and constant availability. This is becoming table stakes in modern commerce. Firms embraced automation because it lowered their cost to serve. The unacknowledged trade off was the reduction in the quality of CX and the lack of differentiation automated systems offered. Up until recently digitisation merely replaced human touchpoints with self-service.

Chapter One introduced the possibility that the gap between human delivered CX and machine delivered CX was closing because the latest AI systems are becoming so much better at emulating humans. The most important implication of this is that the Chief Operations Officer now needs to get the Head of Customer involved in the strategic conversation about what next generation customer fulfilment looks like. A focus on how drive consistent, low-cost operations to achieve scale will not succeed where customers have the option to purchase from a rival offering a better experience. Better experiences require empathy. Humans can design, create, and deliver empathy for a price. Now machines are becoming better able to emulate empathy. They are developing a Theory of Mind.

Chapter Two explored the concept of Theory of Mind. As social creatures we learn to infer what and why other humans are doing, feeling, and thinking the things they are. We use this to explain our observations and forecast what others are likely to do next. This also informs us about how to respond to achieve our own ends in social interactions. We even have a Theory of Mind about ourselves. CX practitioners have tried to develop the marketing concept of personas make customers' Theories of Mind more explicit. This is a precursor to developing customer empathy and designing experiences, interactions and touchpoints that are more customer centric. The experiment at the end of Chapter Two showed just how far ChatGPT had progressed with empathy. The LLM was able to construct valid fictional personas for a bank and then develop key touchpoint interactions. It was able to go beyond derivative CX design when prompted to consider remarkable CX examples from other industries. This HCD was not informed by a specific understanding of real customers.

The takeaway for CX professionals is the role software can play in designing CX. In the same way that Dall-E and other generative AI can create new images by recombining parts of existing images, and ChatGPT can create new stories by generating words that fit together in a meaningful way, AI can emulate human centred design for CX. The CX professional's role in this scenario shifts from quantity to quality: making sure the AI has the right information to understand customers, prompting the outputs it responds with

and evaluating how well it has met the brief. The potential to increase productivity is likely to be between five to twenty fold if the author's experience writing this book is generalisable. Ultimately the AI systems used for HCD should become better than humans because they can produce far more in the same time and take account of a deeper and more nuanced understanding of customer motivations. This reduces the minimum efficient cohort size CX designers and strategic marketers can focus on. Potentially the AI can create individual journey maps and touchpoints in real time with enough training and data.

The challenge evident from the opening chapters is the people who are good at empathizing don't also tend to be good with coding. Chapter Three provided the relevant design criteria for good CX. Good CX is simpler, more satisfying, and hopefully surprising to some degree compared to the alternatives. Simple, satisfying, and surprising are styles of CX. Most companies start out with trying to reduce customer friction to deliver simple experience for their customers. Satisfying experiences are grounded in understanding the customers rational requirements to deal with a company. There is another side to these design principles beyond the rational. Pure efficiency is not often beautiful or emotionally meaningful. A more complete customer Theory of Mind considers customers' emotional needs beyond friction reduction and customers' jobs to be done. The style of experiences can lift their emotional quality and the feelings customers subsequently recall. The most recalled experiences are surprising (whether positive or negative).

Automated experiences can deliver better CX when they are designed to utilise customer data to create personalised experiences. Basic personalisation like filling forms out for customers makes their interactions with companies simpler. Tailoring interactions with customers to match their specific motivations, needs and circumstances increases satisfaction because gross objectives like first time resolution or delivering the right size. Tailoring in real time based on large data sets is where AI works even better than humans. It is harder right now for AIs to surprise humans. Generative AI developments suggest this is only a matter of time. For strategic marketers the opportunity is to leverage next gen AI to provide a more personalised experience than rivals in a way that increases market share. Then by retaining a deeper understanding of customers motivations, needs and circumstances, the firm creates a competitive advantage based on a deeper understanding of its customers. Even when rivals can match interaction customisations, they won't know as much about what changes matter. First mover advantage matters for private data acquisition plays when you can exploit the value in the data. Sadly most organisations are more focused on taking cost out.

Chapter Four and Chapter Five covered the two sides of extracting value from the latest improvements in AI. Automation has historically been used for cost

out. As the automation improved and became cheaper to implement, minimum efficient scale reduced. This opened new opportunities and displaced unskilled and low skill workers. The new opportunities tended to create more jobs than they eliminated in the long run. AI based automation continues the trend. The only difference is highly skilled workers with very specific jobs are likely to be displaced. This includes highly paid workers like doctors, lawyers, accountants and even software programmers.

These workers will have to find new roles and need new skills. History suggests many will try to resist allowing AI to be deployed in their companies so they can keep their jobs. This is as futile as trying to hold back the tide. Their focus should be on taking advantage of the new opportunities created by AI automation. The reason for this is new entrant companies will utilise much more labour efficient business models to compete with incumbents and take their customers. The same logic applies at the level of the firm. Instead of resisting the disruptive AI developments bring, find the opportunities. Chapter Five focuses on how to use the next generation of AI to enhance revenues. This is the core take away for Boards and executive leadership teams developing strategy.

Chapter Six provided two tools to work out where AI can be deployed for advantage. The Pareto Principal highlighted how not all causes are of equal importance to productivity, cost, or quality outcomes. Even if the data suggests in practice the output to input distribution is more likely to be 60/20 than 80/20, there is a critical need to work out what really matters. The fractal nature of Pareto distributions suggests remarkable value comes from identifying the 1% of causes that disproportionately create impact. AI is ideal for this task because it can take account of many more variables in analysing cause and effect compared to humans. It just needs enough data. Once the most important output causes are identified the challenge becomes execution. Theory of Constraints provides a structured process to identify bottlenecks in a business AI automation can reduce or eliminate. In combination these tools can be utilised to create a road map for strategic AI implementation in a business. This will allow resources to be optimally deployed because no business can do everything on offer immediately. This road map is where operations, marketing, CX and finance need to collaborate to design the proposed path to an AI-enabled future.

There is no point upending a business and implementing new technology unless a lasting advantage is retained. The AI strategy should not be about driving a return above weighted average cost of capital. Firms should be looking for 5x-20x gains. This shifts the strategic focus to how to retain their AI-enabled market edge. Chapter Seven introduced two concepts related sustaining competitive advantage. The Resource Based Theory of the Firm was rooted in the idea of exploiting a critical market difference that was hard for

rivals to copy. The Dynamic Capabilities View suggested the only real form of sustainable competitive advantage was the ability for firms to evolve over time to match market changes. In the context of exploiting AI capabilities for competitive advantage three critical elements were identified: control of unique data sets, superior customisation and integration and building on distinctive domain expertise. This chapter is probably most relevant to Chief Information Officers in firms because it touches data, information flows and system architecture design. It is not the direct utilisation of external vendors' platforms that will drive business returns. It will be the unique layer of AI functionality the firm overlays to analyse the privileged customer data and deliver a superior customised CX.

Analysis of public data provides a competitive advantage only if rivals are not conducting the same analysis. A firm's differential data sets (by definition) cannot be utilised by rivals. This will make unique customer insight more valuable in the future. It will increase the focus on data privacy, protection, transparency, and transactions. Chapter Eight extended the concept of trust in AI beyond just data management. The AI not only has to work, but it also must do so in a way that doesn't take advantage of customers. Firms who choose to exploit customers will draw regulatory reactions designed to constrain their activities to profiteer. These will range from onerous and largely ineffective like the GDPR to onerous and effective like the compliance requirements for extending credit to consumers or providing medical care. The key takeaway from this chapter is that a system view is required to manage trust effectively as AI deployments continue. This ultimately is a leadership issue in firms. The responsibility will fall on the shoulders of owners, executives, and advisors to companies to get it right. Again history suggests reasons to be pessimistic about self-regulation. Most firms will try to strike a balance between commercial advantage and customer exploitation. Some firms will cheat and then regulators will overreact after the fact. Then entire industries will be compelled to comply with expensive and unnecessary bureaucracy in the name of protecting consumers. The future outlook is one of more AI capability.

The final chapter suggested a range of potential scenarios. These were derived from identifying significant CX trends supported by the major developments in AI currently happening. More powerful hardware, smarter chatbots, and more prevalence of chatbots will contribute to software and machine user interfaces becoming more human. We are less bullish about the prospects of quantum computing becoming a commercial reality given the effectiveness of the current approaches to cloud and supercomputing. Ultimately the experience of devices will become more embedded in human beings as we want even more powerful, responsive, and pervasive experiences.

We believe the future of AI is even deeper collaboration with 'natural' intelligence to build far higher quality consumer experiences than automation currently provides.

Index

- 5 Ps of CX, 16
- 5Cs of CX, 17
- 80/20 rule, 96
- AbsolutData, 93
- Ada Health, 149
- Adidas, 69
- Affectiva, 33
- Alexa, 104
- ALICE, 23
- Alkco, 109
- Amazon, 31, 132
- Amazon AI recruiting tool, 137
- Amazon Echo, 103
- Amazon warehouse automation, 71
- Amazon Web Services, 63
- Amgen, 135
- Ampex, 62
- Appier, 105
- Apple, 67
- Autofill, 49
- autoGPT, 50
- Automation, 70
- awareness stage, 84
- Axiom.ai, 49
- Babylon Health, 113
- Belfius, 73
- Bizbike, 71
- Blenderbot, 23
- Bluevine, 88
- Boeing, 108
- Boost.ai, 92
- Builder.ai, 135
- Bunnings, 140
- Business process re-engineering, 96
- Cambridge Analytica scandal, 137
- CAPFUL, 46
- Carbon Inc, 69
- CBA, 53
- CEBA, 54
- CES, 59
- ChatGPT, 19
- Chorus.ai, 86
- Cisco, 177
- CLEAR, 48
- Clearview AI, 138
- Cloud computing platforms, 170
- Cogito, 110
- consideration stage, 84
- cost reduction, 66
- Craigslist, 58
- CSAT. *See* SAT
- Customer Effort Score, 59
- customer journey, 81
- customer personas, 35
- Customer Satisfaction, 59
- CX Direct Metrics, 59
- CX lens, 76
- CX Outcome Metrics, 60
- CX roadmap, 115
- decision stage, 87
- Delta Airlines, 107
- DISC personality framework, 104
- Discover EU, 72
- Disney Magic Bands, 57
- Domino's Pizza, 110
- Dramatron, 51
- Drum-Buffer-Rope (DBR), 108
- Duchenne smile, 29
- Dynamic Capability View (DCV), 125
- eBay, 31, 133
- Economies of scale, 66
- Economies of scope, 66
- EDGE (from RBV strategy), 122
- ELIZA, 20
- EmoReact, 32
- Estée Lauder, 105
- EVE Online, 31

Exscientia, 64
 Externalities, 67
 Facebook, 87
 Facebook privacy issues, 137
 Facial Action Coding System (FACS), 32
 Febo, 11
 Flickr, 127
 Foyer, 72
 framing effects, 85
 Freenome, 174
 GDPR, 139
 General Electric (GE), 67
 General Motors (GM), 67
 Gong, 86
 Google Cloud Platform, 63
 Google Project Maven, 141
 Grammarly, 50
 GrammarlyGO, 134
 Grand Rounds, 113
 Graphcore, 171
 Hailey, 92
 Harris Corporation, 108
 Harry's Razors, 51
 Healthcare, 111
 Homo Economicus, 31
 Hugging Face, 175
 Human Centred Design (HCD), 37
 Hyundai, 110
 iFood, 72
 IKEA, 86
 Intel, 107
 Intention to purchase surveys, 11
 IVR (Interactive Voice Response), 52
 KFC, 86
 Kmart, 140
 Large Language Models, 13
 Learning curves, 68
 Lego Boost, 55
 Lego Ideas, 54
 Lemonade (and Maya), 21
 Liberty Bell Telecom, 110
 LLM Hallucinations, 163
 Low-cost inputs, 67
 Lyft, 58
 Mafia Offer, 109
 Magic Castle Hotel, 56
 Marriott, 57
 McDonald's, 66
 Meta, 129
 MetLife, 110
 MindMeld, 177
 Mintzberg's 5Ps model of strategy, 118
 Mitsuku, 22
 natural language processing, 19
 NEC, 140
 NeoFace, 140
 Net Promoter Score, 59
 Netflix, 12
 Neuralink, 180
 Nike, 16
 nonverbal communication, 28
 NPS, 59
 Nubank, 149
 OnDeck, 88
 open banking, 131
 Open source AI, 68
 open source AI strategy issue, 124
 OpenAI, 169
 Optibus, 89
 Owkin, 87
 Panasonic, 62
 Pareto Analysis, 96
 Pareto Principle, 96
 Path dependence, 125
 Personalisation, 46
 Pinterest, 88
 Policy choices, 67
 pregnancy prediction scandal, 106
 Procter & Gamble, 107
 PROFIT (from RBV strategy), 120
 Proximus, 73
 public transport, 89
 purchase stage, 88

Quantcast, 86
quantum computing, 171
Qumulo, 63
Rawbank, 73
Receptiviti, 176
recommendation engines, 12, 13
Remploy,, 109
Rent the Runway, 55
repeat stage, 93
Replika, 34
Resource Based View of the firm
(RBV), 119
Robert Bowden Inc, 108
Rytr.me, 94
Salud Digna, 73
SAT, 59
Sephora, 13
Shenzhen Bao'an Airport (SZX), 79
ShipBob, 88
Sicoob, 71
Siemens, 125
Snapchat, 128
Soul Machines, 178
Speechelo, 179
Speechly,, 49
Stitch Fix, 12
StubHub, 58
Studio Pilates, 52
Sumitomo Dainippon Pharma, 64
supercomputers, 169
support stage, 90
TalkLife,, 92
Target, 106
Tata Steel, 107
Technology Acceptance Model
(TAM, 144
telemetry systems, 132
Tesla, 142
Theory of Constraints, 96
Theory of mind, 26
Thomas Cook Airlines, 56
Ticketmaster, 59
Tinder, 58
Toyota, 68
Trane Co, 109
Trunk Club, 61
trust transfer model, 144
Uber, 150
Unemployment, 74
Unified Theory of Acceptance and
Use of Technology (UTAUT)
model, 145
United Airlines, 80
use stage, 88
Vade Nutrition, 88
Varian's Law, 62
Verified Fan, 58
Virgin Atlantic, 80
Virgin Australia, 80
Visit Wales, 56
Viz.ai, 174
VoiceFiller,, 49
Volkswagen, 141
Walmart, 66
Warby Parker, 47
Waymo, 142
Web 2.0, 126
WebchatGPT, 162
WeChat, 99
Whole Foods, 57
Woebot, 20
Xerox, 110
Yellow Class, 72
Zapier, 90

References

- ¹ <https://www.toprankblog.com/2014/10/people-b2b-marketing/>
- ² <https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/b2b-sales-omnichannel-everywhere-every-time>
- ³ <https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/these-eight-charts-show-how-covid-19-has-changed-b2b-sales-forever>
- ⁴ <https://www.deloittedigital.com/us/en/offerings/customer-led-marketing/close-the-expectation-gap-in-b2b.html>
- ⁵ <https://adtechdaily.com/2020/07/31/b2b-brand-study-97-of-b2b-brands-want-to-become-more-human/>
- ⁶ <https://research.netflix.com/research-area/recommendations>
- ⁷ <https://www.univ.ai/blog/how-does-spotifys-recommendation-system-work>
- ⁸ https://en.wikipedia.org/wiki/Stitch_Fix
- ⁹ <https://uxdesign.cc/understanding-the-end-to-end-user-experience-of-the-sephora-virtual-artist-app-product-try-on-d8ae3f8d1fcf>
- ¹⁰ [https://en.wikipedia.org/wiki/Her_\(film\)](https://en.wikipedia.org/wiki/Her_(film))
- ¹¹ <https://www.nike.com/au/nike-by-you>
- ¹² Dew et al 2021 Lean CX: How to Differentiate at Low Cost and Least Risk p5
- ¹³ <https://appuals.com/best-ai-chatbot/>
- ¹⁴ <https://www.conversica.com/wp-content/uploads/2023/01/Chatbot-Experience-Report-1.23.pdf>
- ¹⁵ <https://venturebeat.com/ai/chatgpt-and-llm-based-chatbots-set-to-improve-customer-experience/>
- ¹⁶ <https://en.wikipedia.org/wiki/ELIZA>
- ¹⁷ <https://medium.com/marketing-in-the-age-of-digital/love-at-first-chat-with-lemonades-ai-chatbot-maya-7b4a105824bd>
- ¹⁸ <https://botpenguin.com/top-7-things-you-didnt-know-about-mitsuku/>
- ¹⁹ <https://www.psychologytoday.com/us/blog/beyond-words/201109/is-nonverbal-communication-numbers-game>
- ²⁰ https://en.wikipedia.org/wiki/Homo_economicus
- ²¹ <https://github.com/bnojavan/EmoReact>
- ²² <https://imotions.com/blog/learning/research-fundamentals/facial-action-coding-system/#>
- ²³ https://www.ted.com/talks/rana_el_kaliouby_this_app_knows_how_you_feel_from_the_look_on_your_face
- ²⁴ <https://replika.com/>
- ²⁵ <https://www.youtube.com/watch?v=qblk7-JPB2c&feature=youtu.be>

- ²⁶ <https://chat.openai.com>
- ²⁷ <https://www.getfeedback.com/resources/cx/how-to-provide-more-personalized-customer-experiences/>
- ²⁸ <https://www.mckinsey.com/industries/retail/our-insights/personalizing-the-customer-experience-driving-differentiation-in-retail>
- ²⁹ <https://www.capfeather.global/capfeather-edition/six-ways-to-differentiate-and-emotionally-cut-through?rq=CAPFUL>
- ³⁰ <https://www.clearme.com/how-it-works>
- ³¹ <https://www.aivo.co/blog/customer-service-challenges-in-banks-and-how-to-overcome-them>
- ³² <https://chrome.google.com/webstore/detail/voicefiller-speech-to-text/mgafhkmkdcbkjajcblfijmlpiknfhffl/>
- ³³ <https://www.speechly.com/blog/turn-any-web-form-into-a-voice-form>
- ³⁴ <https://axiom.ai/automate/form-filling.html>
- ³⁵ <https://www.grammarly.com/>
- ³⁶ <https://support.grammarly.com/hc/en-us/articles/115000090871-How-does-Grammarly-work->
- ³⁷ <https://techcrunch.com/2019/10/10/grammarly-raises-90m-at-over-1b-valuation-for-its-ai-based-grammar-and-writing-tools/>
- ³⁸ <https://whatisthebusinessmodelof.com/business-models/how-grammarly-makes-money/>
- ³⁹
- https://www.reddit.com/r/ChatGPT/comments/zqrwnl/is_chatgpt_going_to_kill_grammarly/
- ⁴⁰ <https://deepmind.github.io/dramatron/details.html>
- ⁴¹ <https://www.primandprep.com/harrys-razors-review/>
- ⁴² <https://www.pilates.org.au/find/faq>
- ⁴³ <https://ideas.lego.com/>
- ⁴⁴ <https://www.lego.com/en-au/themes/boost/about>
- ⁴⁵ <https://www.renttherunway.com/>
- ⁴⁶ <https://www.gsb.stanford.edu/insights/power-free-popsicle>
- ⁴⁷ <https://metamandrill.com/vr-tourism/#marriott>
- ⁴⁸ <https://www.goodhousekeeping.com/life/travel/a40699560/disney-magicband-review/>
- ⁴⁹ <https://blog.contactpigeon.com/customer-experience-examples/>
- ⁵⁰ <https://www.statista.com/statistics/655037/global-direct-to-door-meal-kit-service-market-revenue/>
- ⁵¹ <https://www.theverge.com/2021/10/7/22714762/tinder-lyft-partner-match-free-ride>
- ⁵² <https://www.gartner.com/smarterwithgartner/how-to-measure-customer-experience>

- ⁵³ <https://www.genroe.com/blog/customer-experience-measurement/15139>
- ⁵⁴ <https://www.mckinsey.com/capabilities/operations/our-insights/four-ways-to-shape-customer-experience-measurement-for-impact>
- ⁵⁵ <https://www.mckinsey.com/capabilities/growth-marketing-and-sales/our-insights/prediction-the-future-of-cx>
- ⁵⁶ <https://www-2000.ibm.com/partnerworld/bpmarketintelligence/TheAIRevolutioninCXMeasurement.pdf>
- ⁵⁷ <https://www.cxnetwork.com/cx-technology/articles/roi-is-the-biggest-challenge-to-implementing-artificial-intelligence-in-cx>
- ⁵⁸ <https://www.brittandwhit.com/what-is-trunk-club/>
- ⁵⁹ https://en.wikipedia.org/wiki/Varian_Rule
- ⁶⁰ <https://www.ampex.com/ampex-history/>
- ⁶¹ <https://spectrum.ieee.org/the-consumer-electronics-hall-of-fame-panasonic-pv1563-vcr>
- ⁶² <https://qumulo.com/about/>
- ⁶³ <https://pricinginsight.com.au/walmarts-pricing-strategy/#:~:text=Walmart's%20business%20strategy%20is%20built,labor%20costs%20across%20multiple%20stores.>
- ⁶⁴ <https://www.masterclass.com/articles/economy-of-scope-explainedexplained>
- ⁶⁵ <https://beloved-brands.com/apple-case-study/https://qumulo.com/about/>
- ⁶⁶ https://media.bain.com/Images/BB_Making_move_low-cost_countries.pdf
<https://qumulo.com/about/>
- ⁶⁷ https://media.bain.com/Images/BB_Making_move_low-cost_countries.pdf
<https://qumulo.com/about/>
- ⁶⁸ <https://www.agriculture.gov.au/abares/products/insights/snapshot-of-australian-agricultureagriculture>
- ⁶⁹ <https://www.dpi.nsw.gov.au/dpi/climate/climate-change-and-primary-industries/explore-by-industry/forestry>
- ⁷⁰ <https://global.toyota/en/company/vision-and-philosophy/production-system/https://qumulo.com/about/>
- ⁷¹ <https://d3.harvard.edu/platform-rctom/submission/the-future-of-adidas/>
- ⁷² <https://d3.harvard.edu/platform-rctom/submission/adidas-supply-chain-as-competitive-advantage/>
- ⁷³ <https://d3.harvard.edu/platform-rctom/submission/how-digitalization-is-supporting-adidas-expand-the-reach-of-fast-fashion/>
- ⁷⁴ <https://ourworldindata.org/employment-in-agriculture#employment-in-agriculture-over-the-long-run> <https://qumulo.com/about/>
- ⁷⁵ <https://research.aimultiple.com/robotic-process-automation-use-cases/>
<https://qumulo.com/about/>

- ⁷⁶ <https://d3.harvard.edu/platform-rctom/submission/automated-warehousing-systems-at-amazon/>
- ⁷⁷ <https://www.vox.com/recode/2022/9/27/23373588/amazon-warehouse-robots-manipulation-picker-stowerhttps://qumulo.com/about/>
- ⁷⁸ <https://chatlayer.ai/more-conversions-happier-customers-these-9-companies-use-ai-chatbots-successfully/>
- ⁷⁹ <https://www.bain.com/insights/reduce-replace-rethink-transforming-technology-costs/>
- ⁸⁰ <https://www.forbes.com/sites/ciocentral/2019/08/18/how-compliance-and-cost-reduction-are-funding-data-transformation/>
- ⁸¹ <https://www.brookings.edu/blog/up-front/2022/01/19/understanding-the-impact-of-automation-on-workers-jobs-and-wages/>
- ⁸² <https://www.mckinsey.com/featured-insights/future-of-work/five-lessons-from-history-on-ai-automation-and-employmenthttps://qumulo.com/about/>
- ⁸³ <https://www.weforum.org/press/2020/10/recession-and-automation-changes-our-future-of-work-but-there-are-jobs-coming-report-says-52c5162fce52c5162fce>
- ⁸⁴ <https://www.abs.gov.au/statistics/labour/employment-and-unemployment/labour-force-australia/latest-release#data-downloaddownload>
- ⁸⁵ <https://www.bls.gov/news.release/pdf/empsit.pdfpdf>
- ⁸⁶ https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Unemployment_statisticsstatistics
- ⁸⁷ <https://www.aihw.gov.au/reports/australias-welfare/employment-trends>
- ⁸⁸ <https://www.investopedia.com/ask/answers/061515/what-key-difference-between-participation-rate-and-unemployment-rate.aspx>
- ⁸⁹ <https://www.abs.gov.au/statistics/labour/employment-and-unemployment>
- ⁹⁰ <https://www.oecd.org/economy/labour/33826311.pdfpdf>
- ⁹¹ <https://research.stlouisfed.org/publications/economic-synopses/2015/07/31/a-cross-country-comparison-of-labor-force-participation/>
- ⁹² <https://www.investopedia.com/ask/answers/061515/what-key-difference-between-participation-rate-and-unemployment-rate.aspx>
- ⁹³ <https://www.oecd.org/economy/labour/33826311.pdfpdf>
- ⁹⁴ <https://research.stlouisfed.org/publications/economic-synopses/2015/07/31/a-cross-country-comparison-of-labor-force-participation/>
- ⁹⁵ <https://www.oecd.org/economy/labour/33826311.pdfpdf>
- ⁹⁶ <https://hbr.org/2021/11/automation-doesnt-just-create-or-destroy-jobs-it-transforms-themthem>

- ⁹⁷ <https://www.forbes.com/sites/worldeconomicforum/2020/10/26/dont-fear-ai-it-will-lead-to-long-term-job-growth/>
- ⁹⁸ <https://sloanreview.mit.edu/article/will-ai-create-as-many-jobs-as-it-eliminates/>^{https://qumulo.com/about/}
- ⁹⁹ <https://www.zdnet.com/article/3-new-career-categories-being-created-by-artificial-intelligence/>
- ¹⁰⁰ <https://www.bcg.com/publications/2021/impact-of-new-technologies-on-jobs>
- ¹⁰¹ <https://sloanreview.mit.edu/article/to-cut-costs-know-your-customer/>
- ¹⁰² <https://www.internationalairportreview.com/article/162818/digital-transformation-to-better-serve-airport-passengers-a-case-study-of-shenzhen-airport/>
- ¹⁰³ <https://www.internationalairportreview.com/news/163755/emirates-accelerates-digital-innovation-for-passengers/>
- ¹⁰⁴ <https://www.smh.com.au/traveller/reviews-and-advice/airline-weight-reduction-to-save-fuel-the-crazy-ways-airlines-save-weight-on-planes-20180903-h14v1h.html>
- ¹⁰⁵ <https://www.forbes.com/sites/peterbendorsamuel/2020/08/03/digital-transformation-benefits-beyond-cost-reduction/?sh=1b71ea021755>
- ¹⁰⁶ <https://www.retaildive.com/ex/mobilecommercedaily/neiman-marcus-expands-visual-search-to-provide-instant-gratification-to-consumers>
- ¹⁰⁷ <https://www.forbes.com/sites/bernardmarr/2019/12/09/the-10-best-examples-of-how-companies-use-artificial-intelligence-in-practice/>
- ¹⁰⁸ <https://techcrunch.com/2016/12/23/baidu-and-kfcs-new-smart-restaurant-suggests-what-to-order-based-on-your-face/>
- ¹⁰⁹ <https://www.mycustomer.com/customer-experience/engagement/how-ai-is-improving-cx-at-each-stage-of-the-customer-journey>
- ¹¹⁰ <https://www.bigtincan.com/resources/gong-vs-chorus-how-conversational-intelligence-tools-benefit-bigtincans-customers/>
- ¹¹¹ <https://owkin.com/>
- ¹¹² <https://www.facebook.com/business/help/164749007013531>
- ¹¹³ <https://www.starterstory.com/how-to-start-supplements-business>
- ¹¹⁴ <https://www.business.org/finance/loans/best-same-day-business-loan-lenders/>^{https://qumulo.com/about/}
- ¹¹⁵ <https://www.crazyegg.com/blog/bluevine-review/>
- ¹¹⁶ <https://d3.harvard.edu/platform-digit/submission/how-pinterest-uses-machine-learning-to-provide-tailored-recommendations-to-million-of-users-worldwide/>
- ¹¹⁷ <https://techcrunch.com/2022/05/16/optibus-taps-100m-at-a-1-3b-valuation-for-its-ai-based-mass-transit-operations-platform/>

- ¹¹⁸ <https://zapier.com/blog/ai-in-customer-service/>
- ¹¹⁹ <https://techxplore.com/news/2021-07-ai-chatbots-empathy-gap.html>
- ¹²⁰ <https://www.wired.com/story/new-chatbot-tries-artificial-empathy/>
- ¹²¹ https://www.researchgate.net/publication/336579251_Empathic_Response_Generation_in_Chatbots
- ¹²² https://www.researchgate.net/publication/347388083_Empathic_Chatbot_Emotional_Intelligence_for_Mental_Health_Well-being
- ¹²³ <https://www.weforum.org/agenda/2022/04/artificial-empathy-artificial-intelligence/>
- ¹²⁴ <https://techxplore.com/news/2021-07-ai-chatbots-empathy-gap.html>
- ¹²⁵ <https://www.nature.com/articles/s42256-022-00593-2>
- ¹²⁶ <https://www.boost.ai/about>
- ¹²⁷ <https://www.owler.com/reports/absolutdata-holdings--inc./what-is-a-digital-twin-of-a-customer-and-how-will-/1679602342398>
- ¹²⁸ <https://www.scaleo.io/blog/best-ai-tools-for-affiliate-marketing-future-impact-2022/>
- ¹²⁹ <https://asana.com/resources/pareto-principle-80-20-rule>
- ¹³⁰ <https://www.leanproduction.com/theory-of-constraints/>
- ¹³¹ Sobaca 2010 Achieve More With Pareto's Law p4
- ¹³² <https://www.smartinsights.com/marketing-planning/marketing-models/paretos-8020-rule-marketing/>
- ¹³³ CapFeather client work
- ¹³⁴ https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3498097
- ¹³⁵ Maleyeff 2020 Service Science : Analysis and Improvement of Business Processes p59
- ¹³⁶ Kim et al 2017 The Pareto rule for frequently purchased packaged goods Marketing Letters Vol 28 pp491-507
- ¹³⁷ Tanusondjaja 2022 'Examining Pareto Law across department store shoppers' International Journal of Market Research p 147078532211458-. [online] <https://doi.org/10.1177/14707853221145851>
- ¹³⁸ Jum'a and Basheer 2023 'Analysis of Warehouse Value-Added Services Using Pareto as a Quality Tool: A Case Study of Third-Party Logistics Service Provider' Administrative Sciences Vol. 13, Iss. 2, (2023): 51. DOI:10.3390/admsci13020051
- ¹³⁹
- ¹⁴⁰ Koetsier 2018 'Research Shows Amazon Echo Owners Buy 29% More From Amazon' Forbes <https://www.forbes.com/sites/johnkoetsier/2018/05/30/40k-person-study-buying-echo-increases-amazon-purchases-29-especially-cpg-items/?sh=2fa429e3282e>
- ¹⁴¹ Granata and Palumbo 2022 Impact of Artificial Intelligence on Digital Marketing p103-104

- ¹⁴² <https://www.businessinsider.com/the-incredible-story-of-how-target-exposed-a-teen-girls-pregnancy-2012-2>
- ¹⁴³ <https://www.velocityschedulingssystem.com/blog/drum-buffer-rope/>
- ¹⁴⁴ Cox and Schleier 2010 'A Review of Literature on Drum-Buffer-Rope, Buffer Management and Distribution' Chap. 7 in Theory of Constraints Handbook.
- ¹⁴⁵ <https://www.youtube.com/watch?v=vkqDvLbido0>
- ¹⁴⁶ <https://www.eetimes.com/elec-communications-to-acquire-liberty-bell-telecom/>
- ¹⁴⁷ <https://time.com/5610094/cogito-ai-artificial-intelligence/>
- ¹⁴⁸ <https://time.com/5610094/cogito-ai-artificial-intelligence/>
- ¹⁴⁹ <https://www.health.gov.au/about-us/the-australian-health-system>
- ¹⁵⁰ <https://www.healthline.com/health/what-is-health-equality>
- ¹⁵¹ <https://www.health.harvard.edu/blog/is-our-healthcare-system-broken-202107132542>
- ¹⁵² <https://grandrounds.com/>
- ¹⁵³ <https://www.scmp.com/presented/lifestyle/health/topics/asia-2030/article/3094605/how-babylons-revolutionary-app>
- ¹⁵⁴ <https://www.mindtools.com/a1snnzo/mintzbergs-5-ps-of-strategy>
- ¹⁵⁵ <https://www.theturkeyproblem.com/2018/06/16/nassim-talebs-turkey-problem/>
- ¹⁵⁶ <https://www.blueoceanstrategy.com/tools/red-ocean-vs-blue-ocean-strategy/>
- ¹⁵⁷ <https://journals.sagepub.com/doi/full/10.1177/0149206320982650>
- ¹⁵⁸ <https://www.linkedin.com/pulse/nobody-ever-got-fired-buying-ibm-how-emotional-business-david-forder>
- ¹⁵⁹ <https://taprootmusic.com/the-beatles-road-to-success-in-america-why-record-labels-refused-them-and-what-changed-their-minds/>
- ¹⁶⁰ Nelson and Winter (1982) An Evolutionary Theory of Economic Change Harvard University Press.
- ¹⁶¹ Amy Shuen (2008) Web 2.0: A Strategy Guide: Business thinking and strategies behind successful Web 2.0 implementations
- ¹⁶² <https://www.crunchbase.com/organization/flickr>
- ¹⁶³ <https://www.visualcapitalist.com/timeline-looking-back-at-10-years-of-snapchat/>
- ¹⁶⁴ <https://about.fb.com/news/2021/10/facebook-company-is-now-meta/>
- ¹⁶⁵ https://www.changing-transport.org/wp-content/uploads/2020_tool_strategies_to_maintain_ecodriving_benefits_in_the_long_term.pdf
- ¹⁶⁶ <https://www.freightcourse.com/rail-vs-truck/>
- ¹⁶⁷ <https://backlinko.com/amazon-prime-users>

¹⁶⁸ <https://www.npr.org/sections/thetwo-way/2018/04/18/603750056/long-kept-secret-amazon-says-number-of-prime-customers-topped-100-million>

¹⁶⁹ <https://www.businessinsider.com/amazon-prime-customers-spend-more-than-others-2018-10>

¹⁷⁰

https://ebay.q4cdn.com/610426115/files/doc_financials/financials/2018/q4/Exhibit-99.1-ER-eBay-Q4-2018_FINAL.pdf

¹⁷¹ <https://news.stanford.edu/2017/01/25/artificial-intelligence-used-identify-skin-cancer/>

¹⁷² <https://www.healthline.com/health-news/ai-may-be-better-at-detecting-skin-cancer-than-your-derm#2>

¹⁷³ <https://www.zdnet.com/article/grammarly-is-adding-generative-ai-to-its-platform-heres-how-you-can-use-it-to-elevate-your-writing/>

¹⁷⁴ <https://www.builder.ai/blog/a-little-bit-about-ai-and-more-straight-from-the-builders-mouth>

¹⁷⁵ <https://hbr.org/2022/02/what-makes-a-company-successful-at-using-ai>

¹⁷⁶ <https://www.packagingdigest.com/pharmaceutical-packaging/amgen-invests-ai-assisted-packaging-inspection>

¹⁷⁷ https://en.wikipedia.org/wiki/Facebook-Cambridge_Analytica_data_scandal

¹⁷⁸ <https://techcrunch.com/2023/05/10/clearview-ai-another-cnll-gspr-fine/>

¹⁷⁹ <https://www.eetimes.eu/facial-recognition-the-ugly-truth/>

¹⁸⁰ <https://www.theguardian.com/technology/2022/jul/25/bunnings-and-kmart-halt-use-of-facial-recognition-in-stores-as-australian-privacy-watchdog-investigates>

¹⁸¹ <https://www.wired.com/story/3-years-maven-uproar-google-warms-pentagon/>

¹⁸² <https://en.wikipedia.org/wiki/Waymo>

¹⁸³ <https://kpmg.com/au/en/home/insights/2021/03/artificial-intelligence-five-country-study.html>

¹⁸⁴ Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319-340.

¹⁸⁵ Venkatesh, V., & Davis, F. D. (2000). A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Management Science*, 46(2), 186-204.

¹⁸⁶ Wu, J. H., & Wang, S. C. (2005). What drives mobile commerce? An empirical evaluation of the revised technology acceptance model. *Information & Management*, 42(5), 719-729.

¹⁸⁷ Al-Queisi, K. I., Al-Qatanani, M., Al-Refai, R. M., & Al-Tamimi, H. A. (2019). Factors influencing the adoption of mobile banking in Jordan. *International Journal of Bank Marketing*, 37(2), 456-473.

- ¹⁸⁸ Fathema, N., Shannon, D., Ross, M., & Ali, A. (2015). Integrating mobile learning and achievement motivation with an augmented reality simulator for science experiments. *Computers & Education*, 85, 123-133.
- ¹⁸⁹ <https://building.nubank.com.br/matt-swann-artificial-intelligence-applied-to-financial-services/>
- ¹⁹⁰ <https://ada.com/app/>
- ¹⁹¹ Thurner 2020 The Influence Factors of the patients' usage intention of AI-Based preliminary diagnosis tools: the case study of ADA
- ¹⁹² <https://www.bruegel.org/blog-post/dark-side-artificial-intelligence-manipulation-human-behaviour>
- ¹⁹³ <https://www.bruegel.org/blog-post/dark-side-artificial-intelligence-manipulation-human-behaviour>
- ¹⁹⁴ <https://www.pnas.org/doi/10.1073/pnas.2016921117>
- ¹⁹⁵ Gartner. (2017). Gartner Says By 2020, Artificial Intelligence Will Create More Jobs Than It Eliminates. Retrieved from <https://www.gartner.com/en/newsroom/press-releases/2017-12-13-gartner-says-by-2020-artificial-intelligence-will-create-more-jobs-than-it-eliminates>
- ¹⁹⁶ World Economic Forum. (2018). The Future of Jobs Report 2018. Retrieved from <https://www.weforum.org/reports/the-future-of-jobs-report-2018>
- ¹⁹⁷ See trends.xlsx for 8 websites
- ¹⁹⁸ See trends.xlsx for 6 websites
- ¹⁹⁹ <https://www.verywellmind.com/what-is-episodic-memory-2795173>
- ²⁰⁰ <https://www.cambridge.org/core/books/cambridge-handbook-of-animal-cognition/memory-and-recall-overview/5BD5F0165DF374B0FF361D1147CE2F0A>
- ²⁰¹ <https://journals.sagepub.com/doi/full/10.1177/13623613211030772>
- ²⁰² <https://www.psychologytoday.com/us/blog/the-superhuman-mind/202201/perception-and-the-unreliability-episodic-memory>
- ²⁰³ <https://www.ft.com/content/4329a987-9256-3059-b36f-1aba9338b800>
- ²⁰⁴ <https://news.microsoft.com/source/features/ai/openai-azure-supercomputer/>
- ²⁰⁵ <https://aibusiness.com/nlp/dcw-23-the-cost-and-resources-needed-to-deploy-generative-ai>
- ²⁰⁶ <https://www.fortunebusinessinsights.com/cloud-computing-market-102697>
- ²⁰⁷ <https://www.futuremarketinsights.com/reports/cloud-computing-market>
- ²⁰⁸ <https://www.graphcore.ai/about>
- ²⁰⁹ <https://www.computing.co.uk/news/4022936/british-ai-chip-start-graphcore-close-adding-usd200-million-funding>

- ²¹⁰ <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/tech-forward/the-current-state-of-quantum-computing-between-hype-and-revolution>
- ²¹¹ <https://www.science.org/content/article/biggest-flipping-challenge-quantum-computing>
- ²¹² <https://www.healthcare-brew.com/stories/2023/05/16/existing-ai-technology-could-save-healthcare-usd360b-annually>
- ²¹³ <https://www.freenome.com/about-f>
- ²¹⁴ <https://www.viz.ai/about>
- ²¹⁵ [https://chat.openai.com/c/9a88edbc-ca9d-4920-ae26-86cc94dcf543#:~:text=from_pretrained\(%22bert%20...%22%0AURL%3A-,https%3A//huggingface.co/,-%5B4%5D%20%22Founded%20Date](https://chat.openai.com/c/9a88edbc-ca9d-4920-ae26-86cc94dcf543#:~:text=from_pretrained(%22bert%20...%22%0AURL%3A-,https%3A//huggingface.co/,-%5B4%5D%20%22Founded%20Date)
- ²¹⁶ <https://openai.com/research/gpt-4>
- ²¹⁷ <https://www.receptiviti.com/>
- ²¹⁸ <https://www.mindmeld.com/>
- ²¹⁹ <https://www.mindmeld.com/docs/>
- ²²⁰ <https://www.soulmachines.com/about-soul-machines/>
- ²²¹ <https://speechify.com/blog/ultimate-speechelo-guide-and-reviews/>
- ²²² <https://www.reuters.com/technology/what-does-elon-musks-brain-chip-company-neuralink-do-2022-12-05/>
- ²²³ <https://medium.com/tech-for-good-magazine/march-of-the-cyborgs-8f31acaa2d3e>