

CRIMES AGAINST

MINING PLANNING

Solving The Top 10 Pitfalls

MARK BOWATER



FOREWORD BY GREG LILLEYMAN

Former Rio Tinto and Fortescue Metals Group senior executive

CRIMES AGAINST MINE PLANNING

Solving The Top 10 Pitfalls

Mark Bowater
2022

Copyright 2022 © Mark Bowater

All rights reserved.

No part of this book may be reproduced in any form without the permission in writing from the author. Reviewers may quote brief passages in reviews. The information contained in this book is current at the time of this writing. Although all attempts have been made to verify the information provided in this publication, neither the author nor the publisher assume any responsibility for errors, omissions, or contrary interpretations of the subject matter herein.

Disclaimer

The views expressed are those of the author alone and should not be taken as expert instruction or commands. The reader is responsible for his or her own actions. Adherence to all applicable laws and regulations, including international, federal, state, and local governing professional licensing, business practices, advertising, and all other aspects of doing business in any jurisdiction is the sole responsibility of the purchaser or reader. Neither the author nor the publisher assumes any responsibility or liability whatsoever on the behalf of the purchaser or reader of these materials. All numbers used within both examples and stories within this document are not factually correct and have been created for explanatory purposes.

ISBN: 978-0-6454098-2-6

Acknowledgement

The first person I'd like to thank is someone that without, this book would never have happened. Rob Nixon and I have been good friends for over 10 years and during that time he has written four books in total. Rob coaches accountants and has done so for over twenty years, he is undoubtedly a global thought leader in that space. Rob has told me many times that if you want to make an impact, then you need the credibility that comes from being an author (and in his case a speaker as well). It was Rob who talked me into writing this book and then regularly asked me how it was going over the 14 month journey of writing it. He held me accountable, which was definitely needed considering that I hate writing.

Secondly, my daughter Madeleine was invaluable in producing a final, legible product. Her English is much better than mine (I am an Engineer after all) and so it was a huge relief when she agreed to edit this book and correct my wording, punctuation and grammar. I found completing the first edit of the document extremely laborious and time consuming. Knowing that I didn't have to get too particular about the wording because Maddy was going to fix it up behind me, allowed me to soldier on with editing when at many times it was the last thing I wanted to do.

And finally, but most importantly, I'd like to thank my wife Monica and children Jarryd, Madeleine, Oliver and Kristina. As of 2022, I've now been in business for myself for 23 years. Monica has been an awesome support through that journey, even though there have been plenty of times she has rolled her eyes at another one of my "hair-brain" business ideas. The mining industry is not known for being family friendly and so my family have all had to put up with me spending a lot of time away from home. While I don't call myself an entrepreneur, the quote below helps me continue on this journey and I thank my family for allowing me to do that.

Entrepreneurship is living a few years of your life like most people won't so you can spend the rest of your life like most people can't

- Warren G. Tracy's student

*Trying to predict the future is like trying to drive down a country road at night
with no lights while looking out the back window*

- *Peter Drucker*

*Sound familiar? Mine planning is effectively trying to predict
the future!*

Contents

Foreword VII

About The Author IX

Introduction..... 1

Let Me Be Honest 2

Why Did I Write This Book? 3

Legacy 3

So What Is This Book About? 5

Who Is This Book For? 7

The Structure Of This Book 9

Terminology..... 11

Scheduling versus Planning 12

Schedule Periods 13

Float and Lag..... 13

Dependencies and Interdependencies..... 14

Schedule Components 15

Why Do We Schedule?..... 19

Mine Planning Spectrum..... 23

Life of Mine Plans..... 25

Five Year Plans 26

Annual Plans 26

Quarterly Plan 27

Weekly Plan 27

CRIME 1: The Root of all Evil - Deterministic Scheduling 29

What Is The Problem?..... 30

Why Is Deterministic Scheduling A Problem? 32

The Solution To Deterministic Scheduling 39

CRIME 2: Slippage - What Is Slippage?? 51

What Is The Problem?..... 52

Why Is Slippage A Problem?..... 53

What Is The Solution To Schedule Slippage?..... 62

CRIME 3: Great Mine Plan – What Do We Do With It Now? 67

What Is The Problem?..... 69

Why Is Poor Execution Of The Plan A Problem?..... 72

The Solution To Executing The Plan 73

CRIME 4: You Have KPIs, But A Lack Of Mine Planning KPIs 83

What Is The Problem?..... 84

Why Are Lack Of Mine Planning KPIs A Problem? 87

What Is The Solution?..... 88

CRIME 5: Bullshit In, Bullshit Out..... 97

What Is The Problem?..... 99

Why Is “Bullshit In” A Problem? 104

What Is The Solution?..... 106

CRIME 6: Beware Of The Enablers! 113

What Is The Problem?..... 116

Why Is Lack of Consideration Of Enablers A Problem? 118

What Is The Solution To Enablers? 119

CRIME 7: A Law Unto Themselves 125

What Is The Problem?..... 126

Why Is Not Following The Plan An Issue?..... 128

The Solution To Following The Plan..... 129

CRIME 8: Are You Scheduling The Right Activities? 137

What Is The Problem?..... 138

Why Is Scheduling The Right Activities Critical?..... 140

The Solution To Scheduling The Right Activities 141

CRIME 9: Theory of Constraints 147

What Is Theory of Constraints? 148

What Is The Issue? 150

Why Is Lack Of Use Of TOC An Issue?..... 151

What Is The TOC Solution?..... 152

CRIME 10: Mine Plans Are NOT A Performance Driver! 159

What Is The Problem?..... 160

Why Are Optimistic Plans A Problem? 163

The Solution To Optimistic Planning..... 168

That’s A Wrap! 179

Epilogue - The Broken Stocks Saga..... 187

Recommended Reading (& Bibliography)..... 189

List of Figures

Figure 2 - 1 Task Dependency Types.....	14
Figure 2 - 2 Example Time Usage Model	16
Figure 3 - 1 Mine Planning Spectrum	24
Figure 4 - 1 Electric Shovel Daily Production	34
Figure 4 - 2 Dragline Daily Production.....	35
Figure 4 - 3 Task Interaction – Blasting & Excavation.....	42
Figure 4 - 4 Delay Probability.....	42
Figure 4 - 5 Task Interaction – Waste & Coal	45
Figure 4 - 6 Delay Probability – Waste and Coal	46
Figure 4 - 7 Increased Inventory – Waste and Coal	47
Figure 4 - 8 Reduced Variability.....	47
Figure 5 - 1 Schedule Overview.....	56
Figure 5 - 2 Blast Slippage.....	57
Figure 5 - 3 Five Day Time Buffers.....	59
Figure 6 - 1 Example Gantt Chart	70
Figure 7 - 1 Median or Mean?.....	86
Figure 9 - 1 Total Drill & Blast Inventories.....	123
Figure 10 - 1 Decision Matrix.....	130
Figure 10 - 2 Technical Services (TS) Example	132
Figure 13 - 1 Probability Distribution of Cost Overrun	174
Figure 13 - 2 Required Uplift.....	174

List of Tables

Table 4 - 1 Example Schedule Task Times	37
Table 5 - 1 Schedule Slippage.....	58
Table 5 - 2 Five Day Buffer Schedule Slippage.....	60
Table 5 - 3 Simulation Results	60
Table 6 - 1 Preceding Activity Requirements	76
Table 6 - 2 Following Activity Requirements.....	77
Table 6 - 3 Schedule Management Example.....	78
Table 6 - 4 Top 5 Indicator Example	81
Table 7 - 1 Run Score Example	90
Table 7 - 2 Plan on Plan Change Calculation.....	94

Table 8 - 1 Schedule Error Mapping..... 105

Table 12 - 1 Production Parameters 154

Table 12 - 2 Identifying The Constraint..... 155

Table 13 - 1 Impact Of Continuous Improvement 161

Table 13 - 2 Probability of Finishing Ahead of Plan – Two Tasks..... 166

Table 13 - 3 Probability of Finishing Ahead of Plan – Six Tasks..... 167

Table 14 - 1 Summary of the Crimes 181

Foreword

It is with great pleasure that I endorse *Crimes Against Mine Planning* by Mark Bowater. This book seeks to “call out” examples and themes of poor mine planning and mine management practices that have plagued many of us during our careers in mining.

The mining industry plays a very important role in the Australian economy, with billions of dollars of capital deployed, earning even more billions of dollars of tax and royalty generating revenue. Efficient allocation and use of capital and labour is at the heart of a successful mining operation and good quality mine planning is the key to unlocking the inherent value of an orebody by efficiently planning and scheduling the use of that capital and labour.

Crimes Against Mine Planning seeks to assist the reader in identifying some very common pitfalls and dangers in a typical mine. During my long career, I have seen some pretty average orebodies deliver good returns through some very good mine planning and execution and unfortunately, I’ve also seen the greatest “crime” of all where one of the world’s greatest orebodies struggled to make money through some very poor mine planning practices.

The use of stories in this book are a powerful but simple way to describe what can sometimes be very complex issues. Don’t seek to make things more complex by over-analysing the simple messages within. Having worked with the author and therefore holding some familiarity with a number of the stories within, I can assure you that Mark has faithfully relayed the key messages without over complicating matters.

Given this book is targeted at senior site management and mine planners, I implore you to read this book with an open mind to help protect yourself from being found “guilty” yourself of committing *Crimes Against Mine Planning*.

Greg Lilleyman

Former Rio Tinto and Fortescue Metals Group senior executive

About The Author

I've always found that great, easy-to-read books are built on stories. Stories make for easy reading and present real world opportunities from which to learn. No one wants to be lectured to, but when someone tells a story, well that's enjoyable. So, I have tried to include stories as often as possible throughout this book. Now, I'm by no means saying I have authored a great book, but I have given it a crack. I'll kick that off here and now with my own story, so you know a little bit about the author and where this book comes from. Hopefully, it will contribute to you taking more away from this book, however, it is not compulsory reading, so feel free to jump straight to the first chapter.

After moving to Brisbane, Australia from my birthplace of New Zealand at the age of 12, I settled into public schooling in Brisbane. I attended Ipswich State High School for my two senior years of school and was fortunate to have an awesome Maths teacher, William Bull. Mr Bull spent two years frequently quoting, "If you like this part of maths, then you should think about being an engineer". The problem was, regardless of who I asked, no one could explain to me what engineers did.

The closer I got to finishing high school, the more I had no idea what I wanted to do "when I grew up". In my last year of high school in 1982, my stepfather told me I should do mining engineering. He said that there was only about 40 of them graduating per year across the whole of Australia, so they were snapped up and most of them had jobs locked in by the time they finished third year. To cap it all off, they started on about \$40,000 per year, which was exceptional money for a graduate back in 1982. That sounded like a good plan to me, especially the wages.

So, I applied to study mining engineering at the University of Queensland. I was devastated when I didn't qualify so had to select a different course. I had signed up for electrical engineering at the Queensland Institute of Technology as the backup option but, it didn't take long for three phase electricity and conjugate numbers to confuse the hell out of me. So, after first year I decided electrical engineering wasn't for me and switched to civil engineering.

I loved civil engineering and spent my time trying to decide which field of civil I would move into when I graduated from University. Thus, I promptly forgot all about mining engineering. I graduated with a Bachelor of Civil Engineering in 1987

and scored a job as a graduate engineer with the Main Roads Department on the Gold Coast in Queensland. I was loving life supervising road and bridge construction and completely happy with my career choice until my wife at the time and I visited her parents who lived in Clermont in Central Queensland. My father-in-law worked at Blair Athol Coal Mine and in one of those pivotal moments in life, I was awestruck when he took me on a tour of the mine. I was struck by the size of the equipment and the operation and instantly fell in love with mining. The calling to be a mining engineer returned.

While on the mine tour, I spoke to one of the mining engineers at Blair Athol and asked him how I might get a start in the mining industry. He told me there was a huge shortage of mining engineers in Australia so although jobs advertised might call for mining engineers, I should apply for them anyway, as much of the time they finished up taking a civil engineer. I took his advice and started busily applying for mining engineering positions. It took about a year, but I eventually secured a position as a drill and blast engineer, funnily enough, at Blair Athol Coal Mine. Even more ironically, the vacancy was created when the engineer I had spoken a year earlier took a position at Kaltim Prima Coal in Indonesia.

I spent a couple of years in the drill and blast role before I moved into a dragline engineering role, after which I was eventually promoted to Senior Scheduling Engineer, with the drill and blast, dragline and scheduling functions all reporting to me. The mine scheduling role encompassed planning periods from one week all the way through to life of mine planning.

After four years in the Queensland coal industry, I was itching to get some hard rock mining experience under my belt and, in particular, wanted to work in Western Australia. I subsequently landed a position as Specialist Scheduling Engineer at Paraburdoo Iron Ore Mine, in the Pilbara region of West Australia. The role was in the Mine Planning team and was responsible for long term designs and mine planning from the quarterly plan through to life of the mine. From the day I started in iron ore, I loved it. I found hard rock much more fun and challenging than coal, particularly the importance played by ore quality and cut-off grades. Coal is very black and white (figuratively speaking), whereas hard rock has many shades of grey. It is the geographic locations I have chosen to live in Australia that have dictated my long career in coal, rather than the challenge or work satisfaction.

From the Mine Planning team, I moved into the Production team at Paraburdoo in the role of Grade Control Superintendent. This was undoubtedly the favourite role of my career as a mining employee. Building two product stockpiles every week with eight quality targets in total was both very challenging (see The Story That Started It All chapter) and a lot of fun. During my time, this team was expanded with the addition of a Scheduling Engineer and the Pit Controllers added to the team of Grade Controllers. Our team scheduled from one month down to a shift level and also ensured the schedule was executed correctly. I believe this was the role during which I learned more about scheduling than any other role I have carried out and I find myself wondering if that was because the role involved both planning and execution, as distinct to most planning roles where you complete a plan and then hand it off to a production team for execution.

Following Grade Control, I spent six months in the Production Superintendent role at Paraburdoo leading a team of 72 Operators and 4 Supervisor, before deciding it was time for a move. Having our first child while we were living 4,000 kilometres from our parents and extended families highlighted to us how important those connections are. So as much as I loved working in iron ore and appreciated the amount of learning, it was time to head back home to Queensland. I secured a position as Technical Services Manager for Thiess Contractors at Burton Coal Mine in Central Queensland. Another great role, full of learning opportunities.

What was so good about this role? Burton was a brand new mine, clearing in preparation for mining had not even begun when I commenced in the role. It was a fantastic opportunity to design everything from the ground up and be involved in every decision as to how the mine was going to operate. Including, for example:

- The mining equipment we were going to use;
- The organizational structure of the Technical Services team;
- The software that the Technical Services team would use;
- The mine planning and other technical systems that we would put in place;
- How Technical Services would work with all of the other departments on site; and,
- Last, but not least, I selected the personnel for my entire Technical Services team. I got to choose everyone who was on the team, whereas in most roles you inherit the people that are already there.

Burton was undoubtedly the best and most efficient mine I worked at and this started with the General Manager, who entrusted his managers to do their jobs without micro-managing them. The Production Manager and I agreed very quickly on exactly how things were going to work at the mine site. I fundamentally believe that the Technical Services team is a service provider to the Production team, it's right there in the name of the department – "Services". Technical Services exist purely to be the people with the technical abilities necessary, such as mine design and scheduling, to assist the Production team in fulfilling their requirements. For more on how we worked together, see the story in Chapter 10.

After four and a half years of living in a camp away from home during the week and with young kids at the time, I made a lifestyle choice to leave Thiess and start my own mining consulting company. This choice was ironic because, in my time in the mining industry as an employee, I always had a pretty low opinion of mining consultants. But, from day one, I loved it. It was the best thing I ever did and I should have made the switch much earlier in my career.

The thing I loved most about mining consulting was the variety of work and the high rate of learning that came from working at multiple mine sites. Every mine site is different, they all have their idiosyncrasies and different ways of doing things such as mine planning, reconciliations, design or interaction with the Production team.

From day one, nothing was boring and the rate of learning was greater than spending time at a specific mine site. Over the following thirteen years, I grew my mining consulting company - Echelon Mining Services - from just me, as a single engineer, to a business with thirty consultants and offices in Brisbane and Mackay. However, after thirteen years it was time for a change so I decided to sell Echelon and, from there, actually left the mining industry for about four years. I purchased and managed a locksmith business, which I still own today. During those four years though, I spent my time wondering whether I would let my 25 years of mining experience disappear, never to be seen or used again.

After four years in the wilderness, I decided that my mining days weren't over so I started back in the industry again and here I am today using my thirty years of scheduling experience to improve the issues I see in mine planning.

1

Introduction

Let Me Be Honest

I want to be very clear about this before you get too far into reading this book. I'm not an expert mine scheduler, I've definitely come across people who are better at scheduling than me. I'm not the world's best mine designer, I've worked with people who are better designers. I certainly don't claim to be the smartest person I've met. I'm also not widely experienced in terms of working around the globe in a range of open cut and underground mines. I've met and worked with plenty of mine consultants who's experience lends itself to writing a book on mine planning more so than my range of experience does.

But, I have worked in the open cut mining industry within Australia for over thirty years and have had a continual involvement with mine planning in various forms over that entire time. Mine planning has always interested and intrigued me because of its contradictory nature. Mine plans are generally put on a pedestal at most mines, a sacred set of rules that must be followed, but then there is an underlying belief by many that they are not worth the paper they are written on.

I thrive on change, I believe that if you're standing still, you're actually going backwards. And I love to challenge paradigms and there are certainly plenty of those in the mine planning space.

I commonly hear people say that they don't have time to do something, well we all have exactly the same amount of time, 24 hours in a day, 168 hours in a week. I'm not a believer of "I don't have time to do that", I think it is more a case of "that is not a priority for me right now". So, while others may be better qualified to write a book on mine planning, I made it a priority, as I want to start the journey towards change within mine planning.

I've noticed many flaws in mine planning over the years and held many discussions with other miners about the problems with mine planning. So initially I set out to highlight some of those flaws in the hope that it might help others to avoid them. But I have always believed in the approach of "don't come to me with problems, bring me solutions". So I have included for each mine planning crime my thoughts on some potential solutions.

I don't propose that they are necessarily the best solutions, or the only solutions, but I have provided them to get you, the reader, thinking about possible alternatives that

may be better suited to your mine site. If the one thing I can achieve out of writing this book is starting a conversation on what needs to change in mine planning, then I will feel I have succeeded.

I've tried to facilitate those conversations previously through a group I started on LinkedIn called **MinErs Digs**, but with only limited success, so if you have any suggestions as to how we can facilitate global discussion on improving mine planning – I'm all ears and open to suggestions on how I can help make this happen.

Why Did I Write This Book?

I've worked in the mining industry since 1989, that's 33 years at the time of writing. In that time, I have watched the mining industry go around in circles. I have observed mine planning working poorly and seen very poor integration between the Mine Planning and Production teams, resulting in poor execution of the plan. After 33 years spent what I would call “working in the industry”, I now would like to “work on the industry”, so instead of continuing to schedule, I feel my time is better invested in passing on the lessons I have learned.

Legacy

There is an awesome book I recommend reading called *Legacy* by James Kerr. It is the story of the culture built within the New Zealand All Blacks rugby union team. The basic premise of the culture is that you will leave the All Black jersey in a better place than you found it. This culture is key to the All Blacks being one of the most dominant international teams in the history of sport. In 115 years of playing the game internationally, their win rate is over 80%.

This book inspired me and I have unashamedly borrowed their legacy. When I think about my time in the industry, I would like to leave mine planning in a better place than I found it thirty years ago. I don't believe that overall, mine planning has improved in those thirty years.

When I first started mine scheduling back in 1990, it was a very manual process using spreadsheets and AutoCad, so each schedule took a long time to produce.

However, the manual effort required to schedule ensured that it was generally accurate and that the mine planner had a very good understanding of the plan. Now, the process is totally digitized and takes less labour time, so we churn out schedules more frequently.

Further, because of dramatic improvements in scheduling software over the last 10-15 years, we have added a lot more to our schedules. We typically schedule more equipment, we might carry out destination scheduling as well, we frequently run truck haulage models so we know how many trucks we need and we run optimizer software, so we produce plans that have been maximised for whatever is chosen to be important, such as tonnes or profit margin.

We now spend less time on measuring and inputting the data we schedule on and less time actually scheduling the activities. Mine schedules are definitely far more of a “black box” than they have ever been before. We expect too much of our mine planners and because we believe the mine schedule is accurate and have faith in it, we keep adding things to it. Our faith is a blind faith and I believe we need to stop adding more functionality to our plans and go back to basics, let’s get the simple things right.

Further, we now have less of an understanding of the plan itself due to more time being spent on “add-ons” such as haulage models. Therefore, the core of the plans is not necessarily as correct as it used to be. We are substituting quantity for quality and I believe we rely too much on the software. As we’ve added more functionality, activities and equipment to our schedules, they’ve ballooned in their “busy-ness”. This has resulted in it being significantly more difficult to communicate the plan well. What’s important gets lost in amongst everything else and those charged with executing the plan don’t know what is important and what isn’t.

Along with adding complexity to our schedules, software providers have also enhanced the presentation capabilities of mine scheduling software. For example, many schedule software programs now include animation functionality, so it is possible to watch an animation of the schedule sequence. As more functionality has been added to scheduling software, there has been more to tell in the story and thus the story has become harder to tell. However, we haven’t necessarily improved reporting functionality to cater for that increase in complexity.

So What Is This Book About?

I believe mine planning is misunderstood. Mining companies routinely churn out a range of plans varying from shift or daily plans through monthly plans, quarterly plans, annual plans and 5-year plans all the way to life of mine plans. Planners are frequently under time pressures and it is very common for them to be trying to achieve ambitious schedule targets. These two issues often result in a process of “scheduling by rote”, doing what has always been done and quickly getting a plan out that works.

Often the various inputs to the plan from other teams may be running late, leading to the mine plan being late. So by the time the mine planner finds a plan that works, they have a backlog of other tasks waiting to be finished and it becomes a case of “I got that schedule to work, now I have other things to do”, so they move on to those other tasks.

There is a need for mine sites to sit back and think about all the mine plans that are created and what is the purpose of each of those plans. This book will deal with the purpose of those plans. It also confronts the biggest issues I have observed in mine scheduling in my thirty plus years in the industry. It discusses what the issues are, why they occur and provides suggested solutions to stop their occurrence.

Every site is different, so how significant each of these issues are will vary from mine to mine. I have structured the chapters in order of the severity of their impact on mine planning and, in particular, the performance of the mine site. The issues in early chapters are those that I believe are more prevalent throughout the industry and have a significantly larger impact on mine sites than those in the later chapters. This book tackles head-on why mine planning is not working at the majority of mine sites and it provides a path out of the quagmire we have gotten ourselves into.

This book contains ten chapters specifically covering a range of what I would describe as “criminal activities” that I have seen in Australian mines, as they relate to mine planning. I originally intended for the chapters to be organised in order of their impact (my perceptions anyway), from those having the largest impact to those having the lowest impact. However, that intention didn’t quite work, as one of the most significant issues, that of optimistic plans, incorporates discussion pertaining

to many of the other scheduling crimes as well, so it was more logical to place this chapter last.

So effectively the book is structured now with the two crimes that I believe stand head and shoulders above the other crimes, as the two book ends of the crimes – Chapter 4 (Deterministic Planning) and Chapter 13 (Optimistic Planning).

This is a completely subjective ranking, as every mine is different and what may be a substantial problem at one mine may be of no significance at another. So don't get too concerned about the order of the chapters, don't waste time reading this book and thinking how I've got the level of impacts all wrong. Instead, just bear in mind that the order I have the chapters in is typically representative of my perception of the size of the problem in the Australian mining industry.

Mine planning is, without a doubt, one of the most critical activities carried out at any mining operation, but one of the least understood and one that creates the most arguments and controversy. This book aims to clarify the mine planning myths and unlock what effective mine planning looks like and how to incorporate it at your mine.

This book is written based on my experience in the Australian open cut mining industry, so the problems and solutions relate to Australian open cut mining. This leads to two fairly reasonable questions:

1. Does this book relate to mining elsewhere in the world?
2. Does this book relate to underground mining?

For both questions, the answer is that to some extent, absolutely! As far as I'm concerned, mining is mining and mine planning is mine planning, regardless of whether it is open cut or underground and regardless of which country the mine site is located in. This book is more about the processes than the tools, so I've tried to make it relevant for mines who have no scheduling software or tools and might still be using spreadsheets, right through to those with extremely sophisticated and integrated systems and tools.

All mines involve the completion of a series of tasks to process ore ready for sale to customers. Those tasks could be carried out either in parallel or in series to other tasks and will be dependent on other tasks having been completed first. While there are undoubtedly significant differences between open cut and underground mining,

mine scheduling is ultimately about sequencing those tasks, so the problems and intended outcomes are essentially the same.

And if the issue discussed in a particular chapter is not a problem at your mine site now, there's nothing to say it won't be a problem next quarter or next year. You might be just one change of Manager away from that issue becoming a problem at your site. So I recommend reading every chapter with an open mind and thinking about what is stopping that issue from happening at your mine site.

Who Is This Book For?

It is important that I reiterate that this book is based on my experience and that experience is limited to Australian open cut mines, primarily open cut coal mines. Therefore, the issues and solutions I talk about here are definitely inherent to Australian mining, but mining is a global industry with some very large global players such as Rio Tinto, BHP and Glencore. The prevalence of mining companies with multiple mine sites, often spread across multiple countries, has led to greater uniformity across mines globally. So, while the exact details of the issue or the scale of it might be different from one mine to the next, or one country to the next, I believe it is likely that these issues do exist across the globe.

This book is squarely aimed at mining company management level, that is the General Manager on site, Production Manager, and Technical Services Manager (or Mine Planning Manager). These personnel are generally based on site, but sometimes the Planning Manager roles can also be located at a head office. This book is aimed at these personnel because most of the issues discussed in this book are either created at that level, or the systems and processes that need to change have to be actioned and authorized at that level.

Change is a very challenging process and some of the suggested changes in this book will require significant cultural and process redesign. They will not be easy and they will need a champion at a management level to function adequately. I have observed first hand on numerous occasions how difficult changes are to embed; that is why so many mining companies have change management processes in place.

About five years into my consulting career, I carried out a lot of consulting at a small coal mine over a number of years. By ‘a small coal mine’, I mean it only had one dragline. I worked closely with the operators, supervisors and management team to improve that dragline. Over a period of 12 months, we introduced a number of changes throughout the design, operation and monitoring processes which led to significant performance improvements. However, I was then commissioned to work on other issues at the mine site so the focus on the dragline ceased. It was fascinating (and a great learning opportunity) to then observe how, over time, dragline operations naturally returned to exactly how they operated before we introduced the changes, including reduced performance.

There is a natural inertia that exists throughout the world and things have a tendency to return to their previous state. People fundamentally don’t like change and prefer things “how they used to be”. In no way should you think that the changes suggested in this book, and therefore the changes required to mine planning at your mine site, will be easy. Some of these systems, processes and methods have been in place for a long time.

While this book is aimed at management level, my hope is that it is also very useful for other areas in the mining industry. It goes without saying that mine planners should definitely read this book; it will help to provide an understanding of mine planning and, potentially, the background on why issues are occurring at your mine site.

As Mine Planners, Senior Mining Engineers and Planning Superintendents typically carry little authority at mine sites, it is unlikely they will be able to make all the changes they would like to make. In which case, they won’t have the direct authority to orchestrate change and will need to influence change instead. To achieve that, there is a book that I highly recommend reading, it is called ‘Influencer: The Power To Change Anything’ by Kerry Patterson, et al. The book details a comprehensive model encompassing six sources of influence.

And, if you want to influence change, of course I am always here to help, let me know what we can do. For me, this book is central to my efforts in trying to influence industry change, after years of watching what can be best described as ineffective mine planning.

As with mine planners, I sincerely hope that this book is of value to those responsible for the execution of the mine plans. To those on the Production team, this book will give you a better understanding of mine planning and, potentially of greatest use, it will give you the mine planning background to ask the right questions of your mine planners. It will allow you to challenge the current planning processes on site and create the opportunity for the Production team to be able to influence change as well.

Last, but by no means least, those in the financial modelling and reporting functions at mine sites have value to be gained from reading this book. These personnel are one of the many customers of mine planning and, like any customer, should be interested in the processes that go into the product they are being provided.

In my thirty years in the industry, I have had countless discussions with both site and head office financial personnel and there is definitely considerable confusion between engineers and accountants as to the purpose of mine plans and the deliverables that mine plans should produce. In my time in the industry, I have seen many examples of a lack of alignment between the mine planning team and the financial accounting team. This book will give accountants a better understanding of mine planning and, interestingly, I believe that accountants are potentially best placed to influence change.

The Structure Of This Book

The book starts with a number of chapters that provide background context before tackling the ten crimes. Each of the ten chapters discussing the greatest crimes I have observed are divided into three sections. The first section discusses what the crime is, how it plays out and how it looks at mine sites. The second section is all about the why - why the issue is a crime, the effect that it has at mine sites and how significant an issue it is. The third section in each chapter then addresses potential solutions for how we might stop (or prevent) that crime happening at your mine site.

The last chapter, titled “The Wrap”, is basically the equivalent of an Executive Summary, just at the end of the book instead of at the beginning. Why is it at the end? Because I would prefer that you first read the entire book and gain a full understanding of the background behind these crimes as I believe it will position you better to resolve them. Just reading the executive summary will tell you what

the crimes are and provide a summary as to how to solve them, but I don't believe it will fully arm you for change.

2

Terminology

There are a large range of terms used throughout the mine planning space and, unfortunately, some of them are used interchangeably. There are also terms that some people understand and regularly use, but which are not used by others. This chapter provides an explanation of terms used throughout this book. Note, these are terms used throughout the Australian mining industry.

Scheduling versus Planning

Let's get straight into one of the biggest errors in key terminology used throughout the mining industry. People use the terms "mine plan" and "mine schedule" interchangeably, like they are one and the same thing. I even see this in job titles on a regular basis, Mine Planning Engineer and Senior Mine Planning Engineer, or Mine Scheduling Engineer and Senior Mine Scheduling Engineer. I am guilty of the same thing in my time in the industry. I have used the words interchangeably as well, but I think it is time we get clear on the difference between the two.

Following are the definitions I have decided on to distinguish between the two terms and this is the language that I have used throughout this book.

In its simplest form, a mine schedule is the sequence of activities carried out at the mine site to achieve target outcomes. The mine schedule includes a database of quantities of the tasks to be scheduled, along with equipment productivities, a calendar and then some form of sequence path for equipment to carry out those tasks.

To me, a mine plan is best reflected in the word "plan". A plan is thinking about the set of things to be done to allow something to be achieved. So, in the mining industry when I think of a mine plan, I believe it is the set of thinking that is required to ensure something is going to be achieved (which is ore mined and shipped). In that context, I then think of a mine plan as the mine designs, plus a schedule, plus the communication of it all. That creates a plan for how it's going to occur that can be actioned. So, I will use the term mine plan generally, to reflect the overall process and how it will be achieved, but when I'm specifically talking about the sequencing or scheduling component, then I will use the term schedule instead.

Schedule Periods

There are typically three terms used to describe mining schedules – short term schedule, medium term schedule and long term schedule. Now, again, they are different across every mining company and even mines within the same mining company. For me, a short term schedule is normally up to a period of about three months, so it goes from a shift through to a Quarter. Medium term schedules are typically between three months and two years and long term schedules are anywhere from two years through to life of mine.

The key difference between them is that there is only one of these three schedules that is actually implemented - the short term schedule - which I believe is by far the most important schedule at any mining site. To distinguish these very short term plans from the other range of plans that exist, I'm going to refer to the plans that are implemented as Execution Plans throughout this book. The importance of these plans cannot be understated.

The other two schedules, medium term and long term I would define as information and decision-making plans. Their primary purposes are to provide information and to allow decisions to be made.

Float and Lag

Like the terms planning and scheduling are incorrectly used interchangeably, so are the terms float and lag. Float is the amount of time a task can be delayed without delaying the next task. Lag is a time interval consciously introduced into a schedule to place a delay between tasks.

In the mining industry, it is very rare for us to use the term float in the context of its true definition and also for us to calculate floats within our scheduling processes. This is because float is typically used in the consideration of the critical path and therefore how late a particular task can be before it will delay the project as a whole.

Although we might not use the term in the exact context of its definition, lag is something that is very common within the mining industry. Lags are introduced into schedules as a buffer between tasks, so if the first task runs late it won't delay

the start date for the second task. We don't like equipment parking up because it has no tasks to work on, so we specifically introduce lag into our schedules between most activities, it's just that in the mining industry we achieve that via the use of inventories. We carry inventories so that the second task is delayed from the first task by the amount of inventory carried.

Dependencies and Interdependencies

In the mining industry most tasks have dependencies i.e. when one task is dependent on another. Interdependency occurs when two tasks are dependent on each other, but, in reality, the terms dependency and interdependency are used interchangeably.

There are four types of dependencies, as shown in Figure 2 - 1. These are:

- Finish to Start
- Start to Start
- Finish to Finish
- Start to Finish

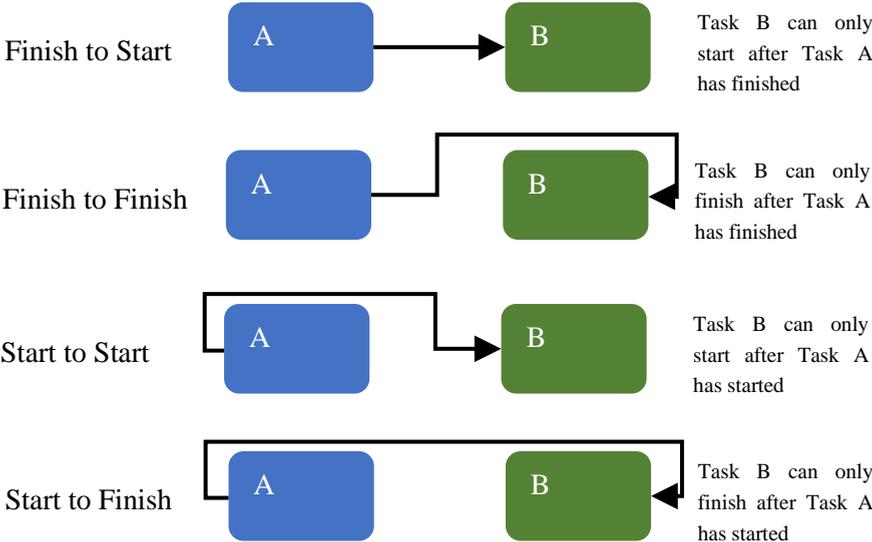


Figure 2 - 1 Task Dependency Types

Finish to Start is by far the most common dependency type and is representative of most dependencies within mine scheduling. For example, the excavator can't start digging the waste until the blast crew has finished blasting it.

Schedule Components

Let's talk about each of the components that go into making up a mine schedule. Within shorter term schedules all equipment will be scheduled so the schedule has more elements to it, whereas in longer term schedules, not necessarily all equipment items are scheduled. Longer term schedules do not require the same level of accuracy and are often simplified to reduce the amount of work required. The following discussion is for execution or short term schedules.

The first item required within mine schedules is what I would describe as the database. The database contains the mining reserves broken down into sub-sets; typically these are on some form of a vertical block basis and also some form of a horizontal bench basis. Effectively, the reserves are broken down into bite sized pieces to allow for more accurate scheduling to be carried out.

Each one of those bite-sized pieces constitutes a record within the database and each record contains numerous fields of information. Within each record, there is typically a significant amount of data, such as:

- geological model results such as raw ore quantities and qualities;
- processed ore quantities and qualities;
- raw waste quantities; and,
- calculations that lead to other quantities which will be used for scheduling purposes, such as drill metres, explosive tonnes and dragline rehandles.

The database provides the core data which is scheduled to arrive at how long each task will take. There is generally some level of transparency in the database to allow an audit of the calculated quantities, meaning that there could be a number of fields required to arrive at some of the calculated values.

Next, you need some form of calendar within your schedule. The calendar will typically have two parts to it. The first is what is typically called a TUM, which is

short for “time usage model”. The purpose of the TUM is ultimately to determine the operating hours that each item of equipment is going to be scheduled. So, we will typically start with calendar hours in a year, then subtract various time losses, such as non-rostered time, idle time, maintenance time and process delays to ultimately arrive at production time (or operating hours). An example of a time usage model layout is shown in Figure 2 - 2.

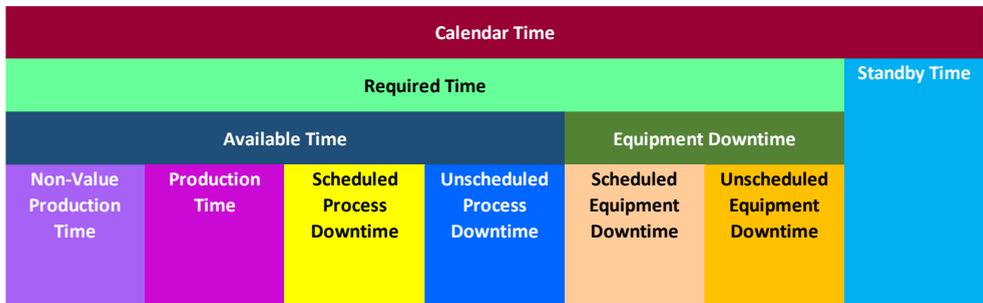


Figure 2 - 2 Example Time Usage Model

Not only does the calendar incorporate a time usage model, it also needs to include specific dates for some time related elements. In a short term schedule it will include dates for days such as public holidays or other time losses. In longer term schedules it will include items such as the following:

- the number of days in the schedule period (say a month);
- the number of public holidays or other non-rostered time in that period; and,
- other elements that are time related, for example wet days in the period, which are often seasonal.

The next element required within the mine schedule is a list of equipment productivities, items such as drill meters per hour, excavated tonnes per hour or ore processing plant tonnes per hour. These productivities are multiplied by the TUM hours to determine the quantities that can be moved within each schedule period.

The remaining part of a mine schedule that is required is the equipment sequence path. For each item of equipment that is going to be scheduled, this is a list (in order) of the blocks or records that will be actioned by that piece of equipment.

Often within the sequence path there is also the opportunity to add delays that fit within the mining sequence, for example deadheading from one block to another.

Finally, there are two schedule components that are not necessarily included in all schedules, but depend on the software that is used for scheduling. The first is output reports, so in some scheduling software you can configure a series of output reports that are produced following completion of the scheduling. These will include mine physicals such as volumes mined and tonnes processed. The second is some form of graphical display. Again, some schedules will allow coordinates to be imported for each of the blocks. This allows block plans to be created which show outputs from the schedule or block parameters such as ore qualities or waste quantities.

3

Why Do We Schedule?

I can't write a book on mine planning without addressing the issue of why we mine plan in the first place and what those mine plans are used for. Particularly since there is such a wide range of mine plans that are carried out, ranging from shift or daily plans, all the way through to life of mine plans and with numerous plans in between. Not only that, but mine plans are carried out with routine frequency, resulting in huge labour time spent on mine planning at every mine site. In fact, I'd hate to even imagine how much time is invested into mine planning each year on a global basis, but needless to say it is many man years.

I think that commonly within the industry, there are questions as to why all of these different types of mine plans are carried out and in addition, there is some misunderstanding as to what the purpose of each mine plan is. I was reminded of this again recently in a conversation with a mining colleague, who commented that he was being challenged a lot about the accuracy of his five year mine plans and was asking me the question "how accurate do five year plans need to be?". We agreed that five year plans did not need a high level of detail and accuracy as, after all, they were never going to be executed. I believe the issue was that this mine had a different reason for running five year plans and, therefore, thought they needed to be very accurate.

Let's start our discussion on why we mine plan by imagining a world where we didn't carry out mine plans and what that world would look like.

We wouldn't have any quantities, so the operations team would tend to develop very good "rules of thumb". For example, the drill completes about 10 holes per shift, there are 120 holes left in this pattern, so that's about 12 shifts of drilling. Or there is about 250 metres of waste left in this strip and we know the excavator digs about 8 meters per shift, so it will take the excavator about 31 shifts to finish this strip.

If there were no mine plans produced, I believe the production personnel (with their rules of thumb), would have a reasonably good handle on what was happening in the very short term, say the next week or so. But get much past a week and it starts to get a bit complex for a couple of reasons. First, the rules of thumb become a bit too inaccurate to work further out and, second, it typically gets too complex at most mine sites. This stems from the multiple equipment items carrying out a range of tasks with numerous dependencies built into the process.

So, without a mine plan, production personnel would probably have a reasonable idea of how much ore was going to be mined in the next week or so, but would have no idea on the quality of that ore, or any success in optimizing the blending. This might be less of an issue somewhere such as black coal mines, where quality is typically less variable and of lower importance. However, in the case of metal mines, where the grade can be extremely variable and there are numerous quality parameters to schedule, production personnel will find it very difficult to plan processing of the ore.

Once we get beyond the immediate short term, it starts to get a bit more complex. Without good timeframes for when the current tasks are going to finish, it is difficult to plan for where equipment will work after the current task because of the high number of dependencies within mining. Mine sites focus strongly on keeping equipment working and it is critical when the current task will finish, allowing for the next task to be planned, so that equipment will continue to work.

Without mine planning, what will be mined and when it will be mined will generally be a guess and there will be no information on product qualities and blending requirements. So, most blending will have to be carried out at the port, based on actual measured qualities once the ore has been processed. It will be difficult to reliably book ships that can be filled and the best way to combat this will again be via rules of thumb, such as booking the ships based on product quantities at the site and port, rather than future planned ore production.

Engaging contractors is a commonly used alternative to capital investment in new equipment and they are also often used to smooth out the volatility in volumes at mining sites. But, if the mine site doesn't have a good understanding of future volumes, then it won't have an understanding of when additional equipment is required so can't effectively plan contractor or future capital requirements. This also means that capital may not be available when required.

If the mine site does not have a good understanding of when contractors are needed then they won't be there to assist when required, due to the lead times before they can actually start on site. But, it gets worse, I would argue that contractors would not be a viable option as there would be way too much variability and uncertainty at the mine site to use contractors. Variability and uncertainty equate to high risk and, for contractors, high risk equals high cost. With contractors you pay for risk, so the

style of contract would have to maximize flexibility, which then comes at a high cost.

But wait, there's more! Here is a list of other issues stemming from not carrying out mine planning and there will be many more I have not included.

- The mine site is not going to know with any accuracy the timing for when infrastructure, such as power lines, will get in the way of mining and will need to be relocated.
- Without mine planning, the mine site won't have an idea of the value lost due to ore sterilization when infrastructure or natural features run through mining areas. For example, what the impacts on future mining quantities will be if a creek is not diverted and that ore is sterilized.
- In Australia (at least) regulatory approvals are now required for many aspects of the mining operation and most of these approvals require comprehensive mine plans.
- Larger mining inventories (blasted waste, ore uncovered, etc.) will need to be carried to cater for inaccuracies within the "rule of thumb" methodology.
- The mine site won't have a good understanding of products they need to order from suppliers, such as explosives. Again, this will lead to the need to carry large supply inventories to cater for the unknown demand.
- It will be impossible to forecast future budget and likely cashflows, leading to substantial variability in earnings and associated volatility in the market valuation of the mining company.
- The mine site won't see critical interactions coming so it is much more likely that there will be increased lost time through deadheading or equipment being idle.
- It won't be possible to align forced equipment downtime events, such as planned maintenance, with other delays when it suits, so during idle time there will be lost opportunities.
- The mine site won't know the areas of the mine with the greatest economics and, therefore, will not necessarily be mining the most profitable ore, leading to underperformance on net present value.
- As the mine won't have a good understanding of when mining areas are going to finish, they won't know when a new mine area will need to start.

Given the long lead times involved in starting new mining areas, the mine won't know when to start the works that are required beforehand.

- The mine site won't be positioned to “make hay while the sun shines” i.e. when ore prices are high, increase production so as to lift profits.

In my thirty years in the industry I've come across numerous production personnel who have very negative views towards mine planning, they wonder why we even plan in the first place and they're prone to ignoring mine plans. The above discussion provides a very good set of reasons why we should mine plan. That discussion concerns the impacts of a complete lack of mine planning, but it also provides some indication of the impacts that may arise if we do mine plan, but the mine plans are not quality plans. So, if we choose to invest in mine planning, it must be because it provides everything we wouldn't have without mine planning.

Mines invest large amounts of time and money into mine planning every year, so why do we plan? What is the purpose of it? I argue that there are two primary reasons why we carry out mine plans. The first is to plan the execution so that it is carried out in an efficient, safe and productive manner. The second is to provide information for decision making or for other informative purposes. Now, let's be very clear about the execution of a plan, there is only one mine plan that is executed and that is the shortest term mine plan that you carry out (depending on the site, this may be daily or weekly). As discussed previously, this plan will be referred to as the **execution plan**.

You can't execute a three month plan or an annual plan, they are likely to be contradictory to the execution plan especially given the mine planning reliability at many mine sites. So why do we measure compliance to plan on anything other than the execution plan? Other length plans are not executed and compliance to plan should only be a measure for execution plans. I believe this alone shows that we fundamentally don't understand how to execute plans. It is one of the greatest scheduling crimes and, therefore, is a chapter in this book.

Mine Planning Spectrum

Many people believe that mine plans operate on a “cascading” basis, where starting from the longest time frame plan (life of mine), each plan feeds into the next shorter

time frame mine plan (for example, the five year plan). So, the life of mine plan should be carried out before the five year plan and provides guidance to that shorter term plan, then, the five year plan provides guidance to the annual plan, etc.

I don't entirely agree with that approach, although there's a good chance you don't agree with me, but hear me out. Here's how I think about it and, therefore, my reasoning. If you turned up at a mine site for the first time and they were carrying out no mine scheduling at all, or they'd just had a fire through their offices and all mine plans had been destroyed, which of the large range of mine plan possibilities would you sit down and carry out first?

That's a no brainer for me, I would complete the Execution Plan - the weekly plan. I can't have Production personnel being aimless and having no guidance for their decisions while I start with a life of mine plan and then proceed on down through five year plan all the way to the execution plan. That approach is pointless when the long term plans have absolutely no impact on the decisions I make in an execution plan anyway. I would then schedule the Quarterly Plan, as it is the decision maker for the Execution Plan, so it answers those short term questions such as when the drill finishes this pattern, where is it drilling next?

After the Quarterly Plan I go to the Life of Mine Plan, I don't believe you can carry out a five year plan if you haven't previously finished a life of mine plan to base it on. It is imperative to understand the life of mine strategy within which a five year plan should be scheduled. Similarly, for the annual plan, it has to fit within the framing derived by carrying out the five year plan. So, I would run my plans in the sequence execution, quarterly, life of mine, five year, annual plan.

Mine plans sit on a spectrum, as shown in Figure 3 - 1, with mine planning cascading down from the life of mine plan to the execution plan, but the execution plan also feeds back into the quarterly plan. It has to, as there are wheels in motion within the execution plan that you're not going to alter in the quarterly plan.



Figure 3 - 1 Mine Planning Spectrum

Given the large range of mine plans that are produced, by far the most common reason for planning is for information and decision-making purposes. So, let's start with the longest term plan and work down. I've set out below the reason for each plan and the primary information that is typically drawn from those plans. Note that mines vary in the time frames for each of their plans, I've chosen the spacing and timeframe for plans that I believe are of most use.

Life of Mine Plans

These are all about providing long term plans for the mine site. They are primarily for strategic purposes, which is the future mining strategy. The following list outlines the types of information and decisions that are outcomes from a life of mine plan:

- What is the remaining life of the mine?
- What is the ore production profile over the life of the mine?
- Are there works required to smooth out that production profile on an annual basis?
- Due to changes in geology, does the ore product need to vary over the mine life?
- What is the waste profile like over that period?
- Are changes required to provide an even spread of the use of equipment?
(For a number of years, I carried out life of mine planning for a large open cut coal mine with numerous draglines. I sequenced the mine on the basis of highest margin coals first. However, this led to some pits finishing significantly earlier and there being insufficient strike length available late in mine life to operate all the draglines. Subsequently, the sequencing needed to be changed so that full operational capacity could be utilised over the life of the mine.)
- What are future capital requirements and the capital profile over time?
- How sensitive is the mine life to a range of parameters such as ore pricing, operating costs and capital costs?
- To instigate or investigate trade-off studies, particularly those with long implementation or pay off periods.

- To provide sufficient notice on very long lead time items (those longer than five years), this includes items such as regulatory approvals around mining lease conversion and creek diversions.
- Determine the life of mine value, therefore facilitating strategic decisions such as whether to divest the mine.

Five Year Plans

These plans are still primarily for strategic purposes, but those decisions that are critical over a shorter timeframe, such as the majority of capital and other items that have lead times less than five years. It is often the case that when these items come into the five year plan, it triggers the required study or other activities to make them happen. Examples are:

1. Equipment purchases.
2. The opening of new mining areas.
3. Infrastructure relocation such as power lines or roads.

Annual Plans

The financial year is very important so an annual plan is important for setting the annual budget. Typically, the annual plan is the first plan that recognizes differences throughout the year, for example, wet season versus dry season.

The annual plan is where you begin the transition from strategic to tactical. This is an important point to note here, as it is important to understand the focus of each plan and to keep it on track by asking whether the purpose of the plan is strategic or tactical. The best way to explain these two terms is to use a military analogy. Strategy is the big picture approach to winning the war - which battles to fight and when. Tactical decisions relate to how you go about winning each of those battles.

In mining, strategic decisions are big picture decisions such as when the best time to open a new pit is. Tactical decisions fit within those strategic decisions, such as, when we mine that new pit, whether we mine in 5 metre or 10 metre depth of benches. To analyse the best bench depth to use, I don't need to run that analysis

out over an entire life of mine schedule to determine the best result, that would be a lot of effort when the annual plan is a long enough time frame to arrive at the right tactical decision. Life of mine and five year plans are much more strategic based, while shorter term plans will involve tactical decisions.

While most annual mine plans are carried out on an individual equipment basis, i.e. all equipment is scheduled, the annual plan is primarily about sequencing decisions involving mining areas or pits. Such as how much production is required in each pit to achieve the annual targets and, therefore, when equipment is going to move from one pit to another. Often it is not critical which piece of equipment it is, but it is more about when the change will be made.

Quarterly Plan

If the weekly plan (the execution plan) is the most important plan on site, then the quarterly plan is the most impactful of all the plans. This is because the quarterly plan is all about setting up the execution plan for success. The quarterly plan is where you are sequencing the completion of activities and where equipment will work next. You are ensuring all pieces of the puzzle fit together so that you are maximising the effective use of equipment, in line with managing inventories and achieving ore targets. If this plan is not representative of reality, it is disastrous for the mine site, as it indicates that the chosen strategy will work. This is the plan that, when wrong, will ultimately lead to equipment parking up or unplanned downtime due to having to relocate to another area to work.

This is the plan that should be based on historical performance rather than an inflated set of assumptions.....

Weekly Plan

This is the execution plan and, I believe, the most important plan by far as this is the one that involves taking actions in the production space to achieve targets. Let's be clear, this is the only plan that is actually implemented, all other plans are for decision or information purposes only. When the quarterly plan is wrong, this is

where it shows up first as you won't be able to get the execution plan to work. By the time you're planning down at the weekly plan level, you are far more constrained and have less flexibility in the decisions you can make. A time frame of a week is insufficient to show sequencing clashes that are coming, that occurs within the quarterly plan.

So, with all these plans to choose from (life of mine through to weekly), if you need to focus effort on improving them, where should you start? If I had to invest in improvement, where would I focus my efforts? I think of it this way, would I rather produce bad information for decisions, or carry out the wrong actions? Personally, I would rather carry out the **right** actions every time. Improving your execution plans is the most important work you can do, as it will cost you money every time you produce a poor execution plan.

4

CRIME 1: The Root of all Evil - Deterministic Scheduling

In my time spent managing my consulting business, Echelon Mining, with over 10 consultant employees, it was important to me that Echelon carried out work that added value to our clients. I believe this is one of the key reasons why I continually focused on my team working at the right level. By this, I primarily mean not working at a level that is much more detailed than required. I commonly came across engineers going into a lot of detail while carrying out long term pit designs, particularly when it came to designing ramps built into the walls. But no one ever builds anything from a long term design, or relocates equipment based on a long term schedule, that's what the short term versions are for.

So, I would walk around the office saying out loud "rough, rough" to anyone within earshot. By this, I meant carry out your work in a rougher (not as accurate) fashion than you currently are. It didn't take long for the team to get the hang of it and to routinely think about whether they were working at the right level of detail.

This phenomenon didn't only relate to design, it was also something that happened in the scheduling area, particularly long term scheduling. Scheduling on a more detailed basis than required is a poor investment of time and there is no better example than the efforts that go into deterministic scheduling. I can't believe the amount of work people put into a plan that is a single snapshot in time and is invariably incorrect before the plan even goes into circulation.

What Is The Problem?

Let's start by zooming out and looking at the typical mine scheduling process used at most mine sites. We're provided a geological model to use for mining designs, from which we export quantities such as volumes, areas, thicknesses and a range of qualities into a scheduling model. We then use a scheduling tool and build in assumptions for the following:

- Equipment operating hours;
- Equipment production rates; and,
- Parameters for turning geological and design volumes into schedule quantities, e.g. converting geological modelled coal volume to product tonnes, prime waste volume to dragline total volume etc.

But, the one area significantly undervalued within our mine scheduling processes, and that I've never seen built into standard mine scheduling, is the inherent variability that exists in mining operations. As an example, let's look at just one activity in the mining process, truck and shovel operations. Here are just some of the variables that occur in this activity:

- Actual dig volume;
- Material density;
- A range of equipment lost time events that are dependent on the activity or other variables, such as:
 - Unscheduled maintenance;
 - Wait on blast;
 - Wait on dozer;
 - Wait on other equipment;
 - Wait on access;
 - Dust;
 - Shovel hang time;
 - Truck queue time;
 - Positioning;
 - Deadheading; and,
 - Idle.
- A range of equipment lost time events that are not dependent on the activity or other variables, but are still variable within themselves, such as:
 - Meal breaks;
 - Shift change;
 - Scheduled maintenance;
 - Refuelling;
 - Crew communications;
 - Wet weather; and,
 - Operator checks.
- (Note the above delays will be different for each equipment item and typically be very different between loading units and trucks.)
- Shovel bucket payload;
- Truck payload;
- Bucket cycle time;

- Truck spot time, travel time and dump time;
- The number of operational trucks;
- Double sided, single sided or top loading; and,
- Face height and face width.

This is just one of numerous activities within a mine schedule, and every other activity, such as drilling, blasting, dragline, coal mining and coal washing, all have a large range of inherent variables. Typical mine scheduling involves consolidating all of those variables into a single assumed production rate multiplied by operating hours that come from a calendar, with single point assumptions for a range of non-operating events.

So, we take a mining operation with hundreds (or more likely thousands) of variables and condense it down to a single snapshot in time that we call a “mine plan” and think that is representative of the mine. We expect an execution team to implement that plan, but then, to exacerbate the issue even further, we invest resources into trying to hold the execution team accountable by measuring “compliance to plan”.

Who are we kidding?!!

This is one of those idiosyncrasies in production that is entirely logical when you sit back and analyse it, but we’re typically too busy in the process of mine design and scheduling to have the time to sit back and appreciate.

Why Is Deterministic Scheduling A Problem?

Is there a bigger waste of time than trying to measure compliance to plan when that plan woefully under-represents the complexity and variability of the mining operation?

Let’s use a very simple example to highlight this. I’m using a simple example so I can calculate the range of outputs using maths, rather than creating a simulation model. However, the real scenario at a mine site is significantly more complex than this example for many reasons, including that mine sites carry inventories,

variabilities are not as simple as normal distributions and mines involve countless interactions between large numbers of dependent activities.

I want to determine the range of total time taken to uncover a block of coal and then rail it to the port. This example involves just one sequence of activities as follows:

1. Drill;
2. Blast;
3. Waste Excavation;
4. Coal Mining;
5. Coal Washing; and,
6. Railing.

There are no inventories, all activities start when the previous activity finishes and they all have the same production parameters. Each activity has an average execution time of 10 days, with that execution time being normally distributed with a standard deviation that is 30% of the mean, so a standard deviation of 3 days.

In this scenario, the average total time for those tasks in sequence would be 60 days and the standard deviation would be 7.3 days. This results in a **41% probability** that the total time will be wrong by more than 10% (so out by +/-6 days) and a **17% probability** of it being wrong by more than 16% (so +/-10 days). It is not uncommon for mine sites to run with inventories of 5-7 days of production, so very simplistically speaking, in that case there is a 40% probability of a schedule issue arising.

An assumed standard deviation of 30% of the mean is potentially on the low side, in analysing a range of real data for draglines and shovels I came up with standard deviations closer to 50% of the mean. In the example above, if we changed the standard deviation from 3 days to 5 days, then there is now a 62% probability of the schedule being wrong by more than 10% and a 41% probability of it being wrong by more than 16%.

But it gets worse!

Have you ever noticed that when a mine schedule is wrong, it's always wrong on the downside? If the schedule was wrong on the upside instead, that would be a nice luxury to have. The result would be that activities would finish early, or the mine site would operate with high inventories. Not such a bad problem to have.

But it never works that way.....

It's always the opposite, instead the activity finishes late, resulting in reduced inventories, equipment being parked up and potential scrambling to mine sufficient product. Why is that? In the earlier discussion within this chapter, I provided example calculations of the total variation, if normal distributions were applied. I kept it very simple and the reason for the use of normal distributions, was that it allowed me to carry out mathematical calculations for the range of total production times.

However, mining production distributions in the real world will rarely be represented by a normal distribution, that assumption is an over-simplification. Production is much more likely to be represented by a left-skewed distribution.

In Figure 4 - 1 and Figure 4 - 2, I have plotted actual daily production results over a nine-month period for an electric shovel in the first figure and a dragline in the second figure. In both charts, the actual results are represented by blue lines. On those graphs, I have also plotted in red a typical normal distribution which matches the assumptions I used in my simplistic calculations earlier in this chapter. That is, the normal distribution has a standard deviation that is 30% of the mean. (Note that for comparative purposes, scheduled downtime has been excluded from the actuals plotted)

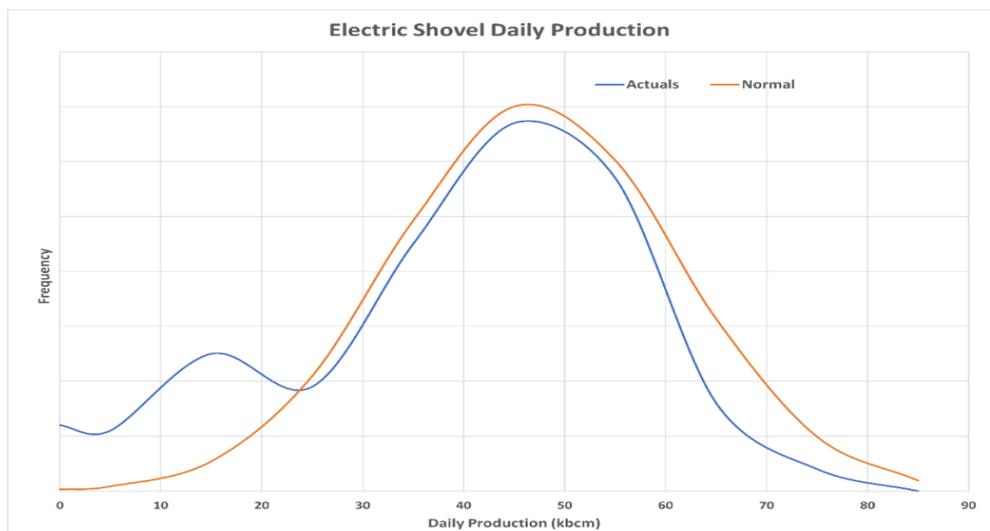


Figure 4 - 1 Electric Shovel Daily Production

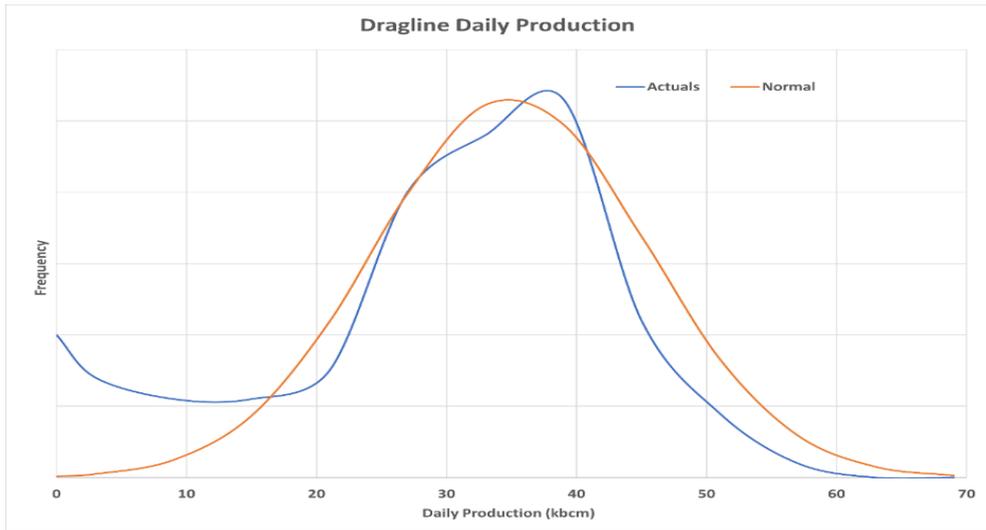


Figure 4 - 2 Dragline Daily Production

The blue lines (actuals) in both charts have a tail to the left (left-skewed). If you think about it, that's entirely logical. The assumed production hours and productivity rates used in the mine schedule will generally have very little upside available in them. During execution, it is generally only possible to create small reductions in delays, such as meal breaks or wait on equipment, or small improvements in productivity rates over those planned. However, it is distinctly possible that delays could be significantly larger than planned, for example, major unplanned maintenance downtime, weather delays or equipment idle time. Similarly, productivity could be dramatically reduced by a poor blast or inclement weather. So, while typically there is very little upside potential on planning assumptions, it is common (and very easy) for significant downside to occur.

The outcome of a left-skewed distribution is periods of reduced productivity, where that productivity reduction is even greater than the scenario where a normal distribution had been built into the scheduling. This leads to activities taking longer than planned, leading to schedule conflicts and, ultimately, failure to meet schedule targets. Unfortunately, more often than not, a mine schedule is going to surprise on the downside rather than the upside.

Of course, this highlights another issue worthy of consideration. That is the issue of **lost opportunities**. Every time equipment is capable of working, but isn't actually working, or is working at a rate lower than capacity, that is a lost opportunity. For example, every time the shovel bucket is in hang mode due to under-trucking, a drill is parked up waiting for bench preparation to finish, or dragline productivity is reduced due to hard digging, that is lost production that can't be regained. Mine sites should focus on minimising those lost opportunities and, hopefully, the introduction of left-skewed distributions into mine scheduling might help to highlight this issue and facilitate improvement.

Mine sites traditionally try and size the fleet capacity for each activity so that it is equivalent to the fleet capacity for each of the other activities. For example, if the mine excavates 50 Mbcm of waste per annum, then the drill fleet will generally have an annual capacity of 50 Mbcm or slightly more. This is called a balanced system.

That all seems very logical and the best way to size your fleet and run your mine, but there is a problem with a balanced operations and the greater the variability, the greater the problem. As always, the easiest and most effective way to show this is by use of an example.

Let's take our theoretical mine from the first section in this chapter, where we tested what happens at a mine with 6 sequential activities that all have a mean completion time of 10 days each. This example is also perfectly balanced, where every activity takes the same time. The mean period from drilling the waste to railing the coal to the port is a total of 60 days. The activities are as follows:

1. Drill;
2. Blast;
3. Waste Excavation;
4. Coal Mining;
5. Coal Washing; and,
6. Railing.

However, each of these activities does not take exactly 10 days, they take 10 days "on average".

In Table 4 - 1 below, I have applied an example set of actual times for each activity which vary around the mean. These are in the column labelled "Time Taken". For

this example, I have specifically set it so that the total time taken of all the activities adds to 60 days, which is the same as the sum of the mean times for each of the activities. The last two columns detail the start and end day for each of the activities.

Activity	Time Taken	Start Day	Finish Day
Drilling	7	0	7
Blasting	12	10	22
Pre-Strip	9	22	31
Coal Mining	8	31	39
Washing	13	40	53
Railing	11	53	64
Total	60		64

Table 4 - 1 Example Schedule Task Times

Now, let’s review the table in detail:

- Drilling starts on Day 0 and only takes 7 days. So, drilling finishes on Day 7, three days ahead of our deterministic schedule (the mean time for the Drilling activity).
- However, because the mine is balanced and we don’t have spare capacity idling waiting for a task, Blasting is likely to be still completing another task elsewhere in the mine. So, Blasting is not available to start until it completes that previous activity, which has an expected time of Day 10. This could obviously be earlier or later than Day 10, but for this example we will assume it finishes at the expected time.
- Blasting will subsequently still start on Day 10 and this activity takes 12 days, so finishes on Day 22, now two days behind schedule.
- Pre-Strip takes 9 days, putting us now one day behind schedule.
- Note that if Pre-Strip completed its prior activity at the expected time, it will be idle for 2 days, waiting for this block to be available to dig.
- Coal Mining takes 8 days, finishing on Day 39 and now putting us one day ahead of plan. However, with a balanced operation, the coal plant is still

washing other coal and is not expected to start on this batch of coal until Day 40.

- Note that if Coal Mining completed its prior activity at the expected time, it will be idle for 1 day, waiting for this block to be available to mine.
- Washing starts on Day 40 and finishes on Day 53, so back 3 days behind schedule again.
- Railing starts on Day 53 and finishes on Day 64, 4 days behind the deterministic schedule of 60 days that was based on mean values for each activity. So, even though the total time for each of the activities was 60 days, the sequential operation actually took 64 days.
- Note that if Railing completed its prior activity at the expected time, it will be idle for 3 days, waiting for this coal to be available to rail.

It is evident in this example that, due to the nature of mining, which is a continuous set of activities, each dependent on the completion of a previous activity, gains were actually lost in the system (unless the next activity was idle and ready to go). However, as soon as a loss occurred, it was retained in the system and passed onto the next activity, therefore accumulating through the system and resulting in lost time of 4 days.

This is one of the key tenements of this book that is completely misunderstood within the mining industry and while we continue to fail to understand this concept, nothing will change.

Gains are lost within the system, while losses are accumulated and pass through the system.

Now, of course, the above is a very simple example and mines are significantly more complex than this, with many other issues at play. But it serves the purpose of highlighting how dependency between activities plays a critical role and this will be further discussed in the following chapter. In this simple example, the first activity of Drilling actually finished early and there was a gain of 3 days. But, for that gain to pass through the system, blasting would have to finish its previous task 3 days early, Pre-Strip would have to finish its previous task 3 days early and so on, all the way through to the last activity of Railing.

If we assume that every activity is normally distributed, then there is a 50% probability of each activity finishing early. The probability of five activities (Blasting through to Railing) all finishing early is 3%. Given each activity is normally distributed, the probability of them all finishing 3 days early is significantly less than 3%. In this simple example, unless we have spare capacity planned within the system, there is less than a 3% probability of the gains passing all the way through the system. If the activity time periods are not normally distributed, but are instead actually skewed distributions, then the probability is even further reduced.

The example above also helps to highlight an issue briefly discussed earlier in this chapter, that is the issue of “**lost opportunities**”. In Table 4 - 1 you can see that Pre-Strip was ready to start on Day 20, but because Blasting was running late, Pre-Strip couldn’t start until Day 22. That is a 2-day lost opportunity that can’t be recouped and there are similar lost opportunities with Coal Mining at 1 day and Railing at 3 days.

The Solution To Deterministic Scheduling

I can’t be any more blunt than this, if you’re not incorporating variability into your mine scheduling, you’re not producing a good mine schedule – it’s as simple as that!

Given the system we are scheduling has a huge number of inherent variables, why are we not incorporating variability and running stochastic models as a standard process for our mine schedules?? We’re never going to create better mine schedules while we continue to run mine schedules on a deterministic basis, that is where they have no variability in the inputs so produce a single output.

So, what do mine sites do to try and mitigate this issue? Mines usually target carrying inventory in front of every activity in the process and typically lots of it. The cost of that inventory across every activity is substantial when you accumulate them across the numerous sequential activities that happen within a mine site. But, even with all that inventory, most mines still consistently run into schedule conflicts that lead to idle equipment or equipment having to relocate.

Further, variability doesn't only play out as an issue that negatively affects mine scheduling, it also negatively affects execution and costs. Production variability leads to the need to carry larger inventories (cost), it leads to lost time opportunities (execution) and also increases the likelihood of a conflict within the mining schedule, requiring the mine to be rescheduled (scheduling).

As posted by a colleague on one of my LinkedIn articles, instead of focussing on larger equipment or other capital investments to lift the mean production levels, maybe mine sites should be focussing on minimising the downside instead and shortening the left side tail of the skewed distribution? This will have the effect of lifting overall production and the added bonus of reducing variability, which allows the possibility of operating with smaller inventories.

It naturally follows then, that I think mine scheduling should be carried out using stochastic scheduling, rather than deterministic scheduling. Just so we're clear on terminology, by stochastic scheduling (also called probabilistic scheduling), I mean that the schedule has ranges built into the completion times for activities, which then require simulation to lead to a range of possible outcomes. Whereas deterministic scheduling is the use of single-point calculations for activity times, with no ranges in the assumptions or outcomes.

Stochastic scheduling is not possible without the right scheduling software. I'm not associated with any software providers so don't have a detailed knowledge of the capabilities of any current mine scheduling software platforms. But, as a potential user of stochastic scheduling software, I would categorise mine scheduling software into three potential capabilities, as follows:

1. Simple – the software displays the probability of an overlap between two activities and, therefore, the probability that one activity will be delayed. However, there is no modification (delay) to any activity start times, even when there is an overlap of the activities.
2. Delayed – when there is an overlap between activities, the software delays the latter activity until the first task is completed. This is more realistic than the Simple option, as this first delay may have follow on effects in later interactions and cause secondary delays and it is critical to understand those.
3. Re-sequenced – in the third option, when there is an overlap between activities, the software reschedules the delayed equipment to another

location so that it does not sit idle. This option is the closest to reality, as most mines would reschedule equipment rather than have it sit idle. However, the ability to reschedule requires the software to have rules to apply and a very smart algorithm to use in determining how to change the schedule.

So what should you do if you don't currently have stochastic capability software and it's not on the radar for your mine to install? At a minimum, you should be thinking about building a manual version of the "Simple" option from above. If you take any interaction between two activities, with knowledge of the activity variabilities, it is possible to determine the probability of one activity being delayed by the other.

You know I love simplistic examples to try and help in explaining, so here's another one. In Figure 4 - 3, we show the variability of two activities, blasting of Block B is planned to finish on Day 15 and it has a standard deviation that is 40% of the task length (not unrealistic from real data I have previously analysed). Excavation is digging Block A and is planned to complete that Block on Day 20 before commencing excavation of Block B (which has just been blasted). It also has a standard deviation of 40% of the task length. (Note I have assumed normal distributions for both activities, but as per previous discussions, it is more likely that both curves will be skewed).

These two activities have a five day lag between them, which is also not uncommon in the mining industry. But, as you can see, there is a large degree of overlap between the two curves. That overlap represents the probability that there will be a schedule conflict, which means that the excavation activity will not be able to start as blasting is not finished.

Figure 4 - 4 shows that the overlap in this case works out to be a 27% probability that excavation will be ready to start before blasting has finished and will therefore have to wait for the blast.

I would suggest that for the most critical interactions in the mine schedule, it would be worth carrying out this exercise to understand the probability of there being a schedule conflict. It would be relatively simple to set up some spreadsheet templates with historical variability distributions and then just use the planned lag from the schedule that is being tested.

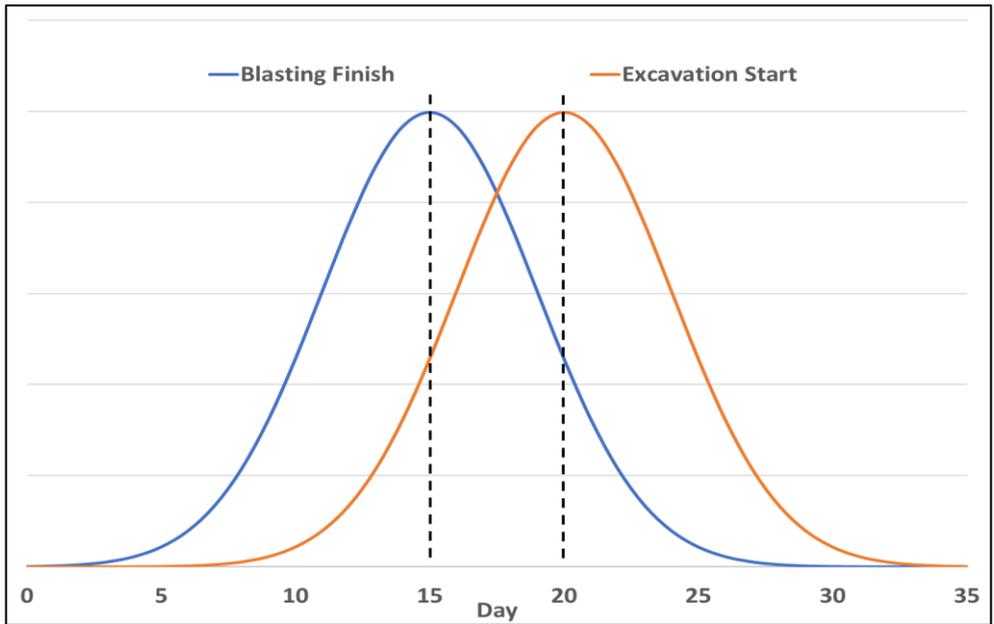


Figure 4 - 3 Task Interaction – Blasting & Excavation

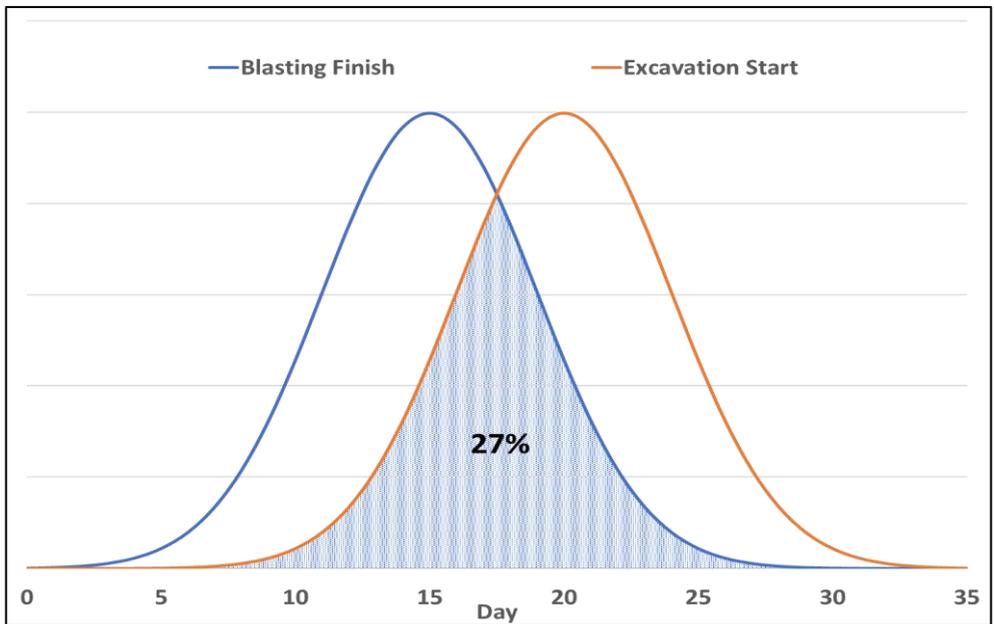


Figure 4 - 4 Delay Probability

With this tool, it is then possible to carry out further analysis to understand how sensitive the interaction is to changes in the schedule. For example, in this scenario, if the lag between activities was only 2 days instead of 5, then the probability of excavation having to wait on blasting increases to 40%. But, if the lag was 8 days instead, then the probability reduces to 16%. If you want the interaction probability to reduce to below 10%, the planned lag has to be 11 days or longer.

It would be a valuable exercise to incorporate this last set of logic into the exercise of creating targets for float between activities when scheduling. Consider what risk of interaction between activities you are prepared to accept. This then allows you to determine the target minimum lag times to use between activities when scheduling.

This raises the issue as to how you decide which interactions in the schedule will likely have the highest probability of an overlap. Here's a method that will provide a very good indication of where to start with your calculations. Divide the lag between two activities by the average standard deviation of those two activities. The lower this number, the higher the likelihood of there being a schedule conflict.

So, there it is, even if you can't source a software tool capable of incorporating variability into your mine schedules, you now have the logic to create a simple tool that allows you to understand the inherent risk in your schedule. The lower the risk, the longer the life of the schedule (before it is wrong) and the greater the likelihood of being able to successfully implement the plan.

As a side note, if you don't have the ability to carry out stochastic scheduling, you should consider how to de-risk a mine plan. There's probably an entire book in discussing how you de-risk all of the components of a mine plan, so this section will discuss just one component. The primary risk inherent in a mine plan is that it is wrong, in which case it is redundant. No mine should be making decisions based on a plan that is wrong nor trying to execute it.

A mine plan can be wrong when it is created, the primary causes being:

- Incorrect geological model or interpretation of it. There are many people significantly more qualified than me in this area so it will not be discussed in this book.

- Errors within the underlying schedule such as numbers typed over formulae, incorrect naming or numbering of blocks, etc. This issue is easily resolved if the planner takes the time to carry out sufficient due diligence on the schedule before starting to plan with it. Unfortunately, planners often have insufficient time prior to scheduling to carry out proper due diligence (see Chapter 8).
- Use of the wrong assumptions in the plan. This is prevalent throughout the industry and by far the biggest culprit is the use of optimistic assumptions. It is usually driven from an executive level within the company and involves setting unrealistic schedule targets or dictating assumptions that are higher than historical performance. Optimistic scheduling is a disease and is further discussed in Chapter 13, as I believe it is the single largest cause of poor scheduling within the mining industry.

Alternatively, the mine plan can become wrong at some time during the life of the plan. If there are no database errors in the plan, then it becomes wrong for one reason – **variability**. Variability happens across all components of the plan, from the scheduled quantities, to equipment productivities, to equipment operating hours.

If you want to de-risk a mine plan so that, once created, it is more likely to be sufficiently correct to be useful, then you have two options:

1. Decrease the variability in those elements which you have control over, which will reduce the spread of time taken for tasks, reducing the probability of overlap between those tasks.
2. Introduce longer lag times between tasks, that is the time between one activity finishing a block and the next activity having to start on that same block.

Let's look at the relative impact those two changes have on a mine plan. To do this, we can use a simple example and the same tool that we used previously. In this example, we are excavating a block of waste (Block B) to uncover coal for mining. The coal mining fleet is excavating another block of coal (Block A), before relocating to mine Block B coal once it has been uncovered. Let's assume the following:

- The variability of both activities is normally distributed.

- The waste excavator is planned to take 10 days to dig Block B and the variability in this activity is represented by a standard deviation that is 40% of the average, so a standard deviation of 4 days.
- The coal mining fleet is planned to finish the excavation of Block A at 5 days after the waste excavator finishes Block B, so there is a lag between activities of 5 days.
- Coal mining of Block A is planned to take 6 days and the variability is represented by a standard deviation that is 40% of the average, so a standard deviation of 2.4 days.

Figure 4 - 5 shows the spread of potential finish times for the waste excavation task and start times for the coal mining task.

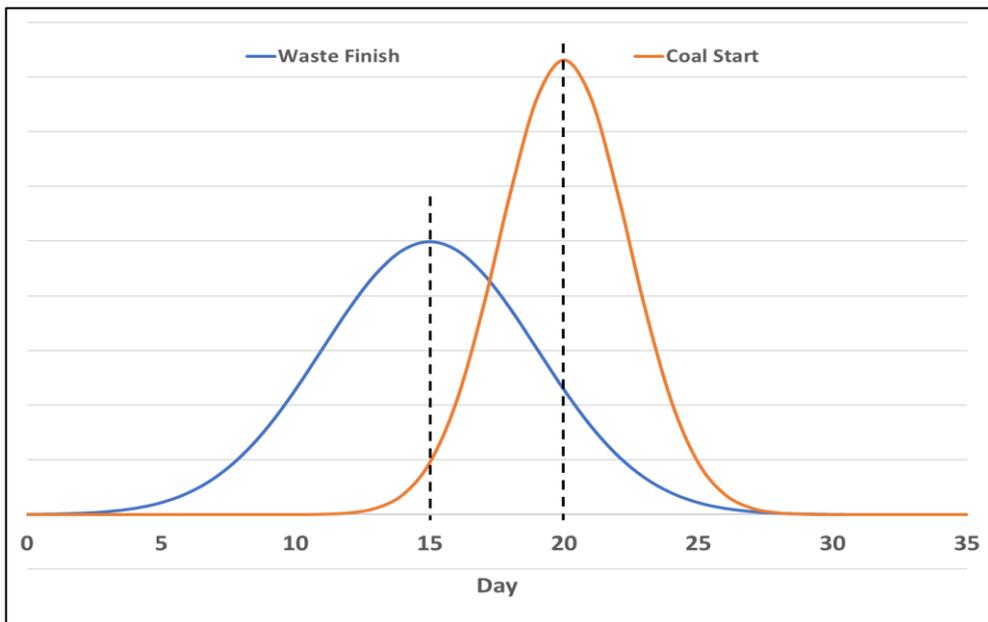


Figure 4 - 5 Task Interaction – Waste & Coal

As shown in Figure 4 - 6, this scenario leads to a 21% probability that there will be an overlap between the two activities, and consequently coal mining would be delayed by waste excavation.

De-risking the plan involves reducing the probability of an overlap, as an overlap will result in either a re-sequencing of activities or coal mining being delayed. But, either way, the mine is no longer operating at its maximum efficiency and the schedule is now wrong.

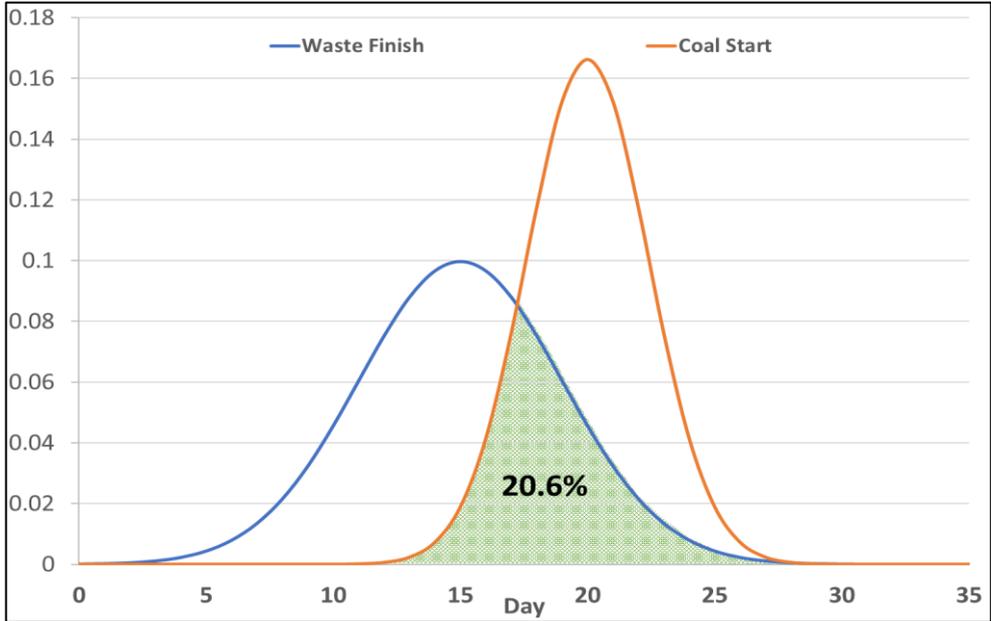


Figure 4 - 6 Delay Probability – Waste and Coal

To reduce the probability of an overlap requires either an increase in inventories or a reduction in activity variability.

Let’s look at the impact of increasing inventories first, which, in this scenario, leads to an increase in the lag between activities. If we increase inventory by 30%, that would take the lag from 5 days to 6.5 days. As shown in Figure 4 - 7, this reduces the probability of an overlap from 21% to 14.7%.

Now let’s look at de-risking by reducing activity variability and use the same percentage change as that used in inventory change. So, we have decreased the waste excavation standard deviation by 30%, from 4 days to 2.8 days, and for coal mining, the standard deviation has reduced from 2.4 days to 1.7 days. Figure 4 - 8

shows the impact on the probability of overlap, which has decreased from 21% to 12.5%.

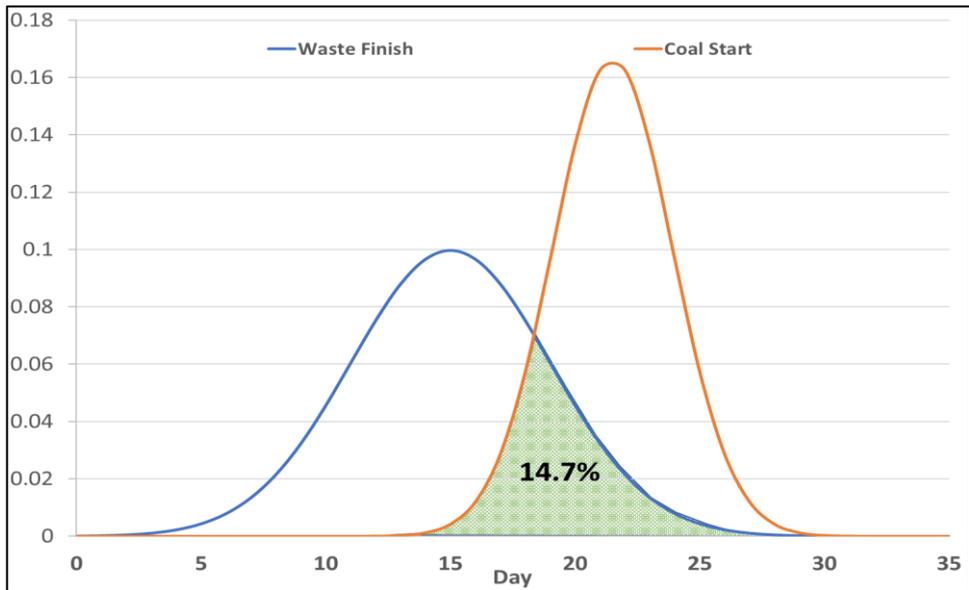


Figure 4 - 7 Increased Inventory – Waste and Coal

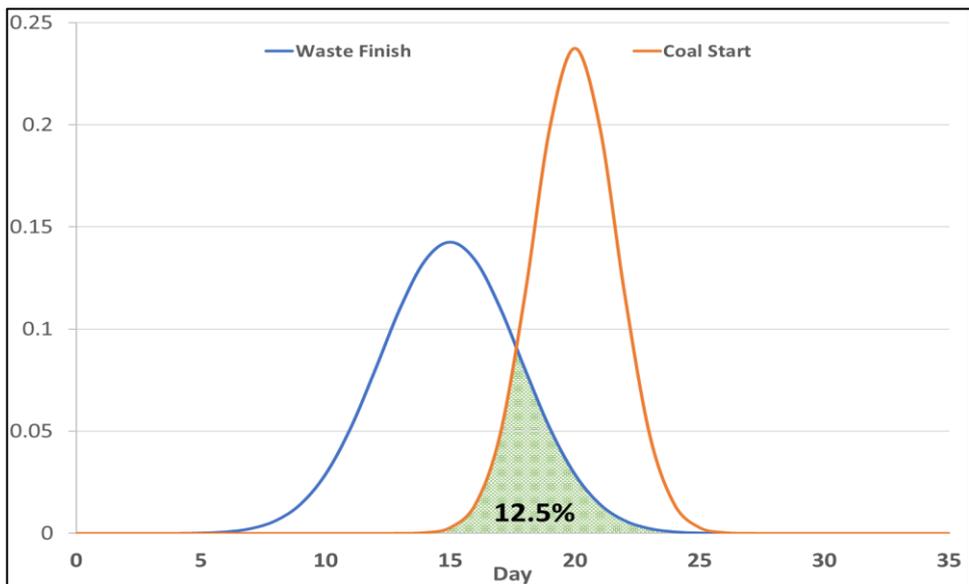


Figure 4 - 8 Reduced Variability

In this scenario, the probability of overlap is more sensitive to a reduction in variability than it is to a change in lag. However, that may not be the case if the parameters were different, for example, significantly longer lags or reduced variabilities.

Now let's look at the financial impacts, starting with changes in inventory. Inventory is effectively cash spent earlier than it needs to be, and the ongoing cost of that money spent early. The cost of debt at the moment is extremely low, however, for most mining companies the consideration is not the cost of borrowing. Mine sites always have a range of capital investment choices, such as purchasing new equipment, developing new mining areas or debottlenecking processing plants. So, money should only be invested in inventory growth if it has a better return on capital than other investment opportunities. Obviously, that is extremely variable across mine sites, but for the sake of this example, let's say that the minimum return on investment required is 40%. Let's also assume truck and shovel excavation costs \$4/bcm to dig, plus an additional \$1/bcm to drill and blast, so a total cost of \$5/bcm for the additional inventory.

Let's assume that this mine has five 600t sized diggers and each digger excavates 9 Mbcm per annum. This mine excavates 45Mbcm per annum of waste, which equates to 125,000 bcm per day. If we want to increase the lag by 1.5 days, then we will need to excavate an additional 187,500bcm of waste at a total unit cost of \$5/bcm. At a holding cost of 40%, this equates to an annual holding cost of \$375,000.

For a mine excavating 45 Mbcm per annum of waste, that is not a huge cost in the scheme of their total budget; however, that is for only one activity and the mine will be carrying inventories for every activity. This is also just the cost of increasing truck and shovel waste by 1.5 days, so is not representative of the annual cost of the total inventory carried, which would be a large value for most mine sites.

But, let's consider what it might cost to reduce variability. That is not a simple calculation, as reducing variability will require a systematic change process. There is every likelihood it will conform to the 80/20 rule, where 80% of the reduction will come from tackling 20% of the issues that cause variability. Let's take mine site culture for example, I've lost count of the number of mines that proudly report a

“record shift” when it happens and make a big deal out of it. However, they never talk about what the production was for the preceding or following shifts.

Boasting about record “one-off” efforts only drives a culture of increasing variability and highlights the complete lack of understanding that mining companies have of the impact of variation. Never yet have I heard a mine site boast of 10 shifts in a row with production between 13,000bcm per shift and 17,000bcm per shift. Never yet have I ever seen a KPI which is a measure of a lack of variability.

There are simple (and free) changes that would start a shift towards reducing variability, instead of increasing it. So, a mine site could spend \$375,000 per annum on additional waste inventory, or it could start a systematic change process aimed at reducing variability, with many of these changes coming at no cost. I know which one I’d be doing.....

5

CRIME 2: Slippage - What Is Slippage??

A colleague of mine, Martyn, and I have had many, wide ranging discussions regarding the inherent problems within mine scheduling. For a while, he was regularly referring to the term “schedule slippage” and it usually involved drawing on a whiteboard. I made a comment to him one day that it was interesting to me that slippage wasn’t really a term I’d come across before in the mining industry. He replied, “that’s right, because I created it!”

Slippage is the delay in the schedule outcome resulting from tasks being dependent on multiple other tasks. Which begs the question, why isn’t slippage something that is acknowledged more often in the industry??

What Is The Problem?

Have you ever noticed that mine plans are much more likely to underperform than overperform? That mine sites are more likely to be behind plan than ahead of plan? Here’s why.

In discussing probabilistic planning with Martyn, he told me it had recently dawned on him that a plan with P50 input assumptions does not result in a plan with P50 outputs. This was not something that I had thought about previously.

To explain what this means, P50 is a probability of 50%, so P50 input assumptions, such as equipment productivities or utilised time, mean there is a 50% probability that the equipment will achieve that assumed production level. Whereas a P50 output means there is a 50% probability that the plan as a whole would be achieved, which is usually assessed by whether ore targets are met.

In writing this book, I like to provide good background discussion and I think there is no better way of getting your point across than using examples for illustrative purposes. This is because these concepts might be obvious to the mine planners who work with them every day, but unless the executive team responsible for decision making and the running of the mine are aware of these concepts and understand them, then we’ve had minimal impact. So this section is written more for the **executive team** than it is for mine planners.

I initially set out to answer the question about P50 in versus P50 out in a short discussion but given my desire to provide purposeful explanation and use telling

examples, it became longer than anticipated. So, I'll start by providing some background discussion and concepts, along with giving you some indication as to which way the answer is likely to fall.....

As I pondered Martyn's statement, I realised that the issue of slippage very closely relates to what I communicated in Chapter 4, in my discussion on how losses are passed through the system but gains are lost. Therefore, this Chapter reiterates much of what I discussed in the previous chapter, but this is a critical issue so it is worth discussing again and communicating in a different way if that helps increase understanding of this issue within the mining industry.

Does P50 in equal P50 out? There is a significant contributor to this question which is rarely discussed in mine planning circles and I don't remember reading about it in any training manuals or academic guides. The contributor is "slippage". Slippage arises when a task has multiple dependencies, that is, the task is dependent on more than one other task finishing before it can start (or be finished).

It is rare in mining schedules that a task is not dependent on at least two previous tasks. For example, if we are going to excavate Block B, the excavator allocated to this task must finish digging the previously scheduled block, Block A, before it's available to dig Block B. This is known as a resource dependency, that is, a task depends on a resource being available. But, for the excavator to dig Block B, it must also be blasted, so the blast crew must have blasted Block B. This is known as a task dependency, another task in the sequence must be finished first.

Why Is Slippage A Problem?

The following discussion shows why multiple dependencies are such an issue. Let's take the example above, in real life we're not going to finish the relevant tasks exactly on time, that only happens in our deterministic plans! Each of these tasks is either going to run early or late. If we've chosen P50 for our input assumptions, then there is the same probability of running early as there is of running late. So, we have four possible scenarios:

1. Digging Block A finishes early (50% probability) and blasting Block B finishes late (50% probability). This option has a 25% probability overall and results in Block B excavation starting late, because it wasn't blasted.
2. Digging Block A finishes early (50% probability) and blasting Block B finishes early (50% probability). This option has a 25% probability overall and results in Block B excavation starting early because it was blasted and the digger was available.
3. Digging Block A finishes late (50% probability) and blasting Block B finishes late (50% probability). This option has a 25% probability overall and results in Block B excavation starting late because it wasn't blasted and the digger wasn't available.
4. Digging Block A finishes late (50% probability) and blasting Block B finishes early (50% probability). This option has a 25% probability overall and results in Block B excavation starting late because the digger wasn't available.

Of the four combinations, only one is ready to start the second task early, the other three all result in a late start. There is a 75% probability of slippage occurring within this schedule, but it gets worse. The fact is, other than for the first task in the sequence (such as drilling the uppermost bench), this is pretty much as good as it gets. At times a task may be dependent on three other tasks rather than two, so, for example, digging Block B may be dependent on blasting Block B, the digger having finished Block A and the mine services crew having established a sump and dewatered the area. In this scenario, we now only have one out of the eight possible combinations that will lead to an early start. If we're planning with P50 inputs, there is now only a 12.5% probability of an early start of Block B excavation and an 87.5% probability of slippage.

What happens if we're not planning with P50 inputs? What if, as many mines are prone to do, we're using P40 inputs instead? Now there is only a 40% probability of each task finishing early and a 60% probability of it finishing late. It turns out that this small difference has a significant impact. Our first example above, which had a 25% probability of starting early, now only has a 16% probability ($40\% \times 40\%$) of an early start. Further, our second example, which had three dependencies, now only has a 6.4% probability of an early start or, expressed in a more compelling manner, a 93.6% probability of starting late.

These are simplistic examples with the aim of communicating the importance and impacts of slippage, however, in them I've ignored the concept of inventories. But inventories have several impacts on the probability of slippage. We use inventories in mining primarily to ensure that equipment can continue to work, given the variabilities we have in mining processes. So, continuing with our first example above, having blasted inventory would mean that the excavator still has Block A blasted inventory in front of it to dig on the date that Block B is scheduled to be blasted. That is, Block B is scheduled to be blasted before Block A excavation is scheduled to finish. Now, at this stage, the maths gets too complex for a simple mining engineer such as me so I'm going to use simulations to assist in further exploration and arrive at a conclusion as to whether P50 in equals P50 out.

Initially I've tried to explain slippage using only mathematics, as that was all I had at the time to use in arriving at solutions. But there came a point where to fully explain the impact of dependencies, the use of simulation software was required. Simulation tools allow the user to build variability into input parameters, allowing scheduling to be carried out on a probabilistic basis, giving me the ability to carry out more comprehensive examples and therefore highlight critical concepts.

I want to highlight why P50 in does not equal P50 out and that in fact, the differential is so large, even I was shocked at the size of the gap. As I've done previously, I am still going to run with a fairly simple case with 6 sequential activities, so each activity must be finished before the following activity can be started. Those activities are:

1. Drill;
2. Blast;
3. Excavate waste;
4. Excavate ore;
5. Process ore; and,
6. Rail ore.

In this example I have assumed the mine is a balanced system, that is the equipment is all sized so that it has the same annual capacity, which means there is no planned idle time stemming from excess capacity. This is typically representative of the mining industry, as one of the strongest drivers in our industry is to not have equipment idle, it must be working at all times.

Because the capacities are balanced, all tasks take the same time period. In this example I have assumed 12 days for each task. All tasks are normally distributed and have a standard deviation that is 30% of the task length, so 3.6 days.

In this example mine, there is only one of each equipment type, so it is fully sequential. The drill can't start one block until it has finished another block, as there is only one drill on site. I've simulated the mine over a long enough period that continual operation exists, so all equipment have tasks to be completed throughout the entire schedule. A snapshot of a portion of the schedule can be seen in Figure 5 - 1.

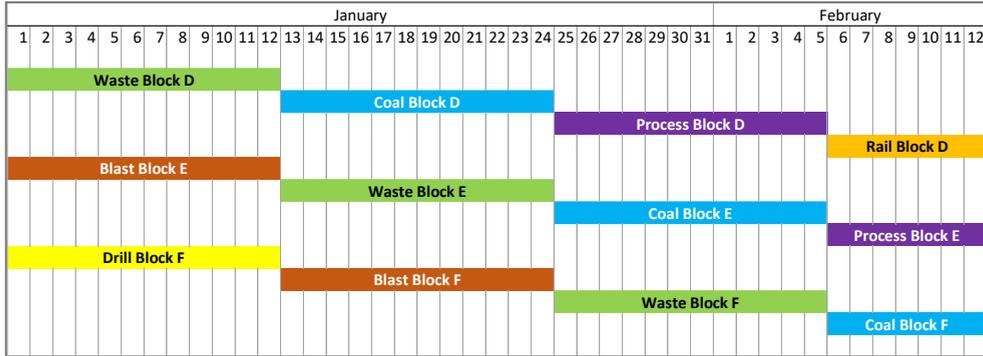


Figure 5 - 1 Schedule Overview

As discussed previously, slippage occurs when tasks have multiple dependencies. For example, Block C can't be blasted until it has been drilled, but there is only one blast crew, which is not available until after they have blasted Block B. So, in this simple example, I have looked at a snapshot of the schedule starting on January 1st, when we commence riling Block B and have scheduled through to April. However, for comparison purposes, the date I am interested in is the date when Block F completes riling. In the deterministic schedule, there are 6 blocks of coal to be riled at a task length of 12 days each and with no gaps between those tasks, that is 72 days and so Block F completes riling on March 13th.

Figure 5 - 2 displays the impacts once we introduce the variability of 3.6 days standard deviation for all of the tasks. Blasting Block E and drilling Block F are at

the start of the schedule (January 1st) and so will have variability around their end date, as shown by the grey distributions. But, these tasks are not dependent on any other tasks and therefore have a 50% probability of finishing on January 13. Blasting Block F, however, is dependent on the blast crew having blasted Block E and the drill having finished Block F and both activities could finish early or late. But as discussed earlier in this chapter, in 75% of the cases a delay will be passed through the system to the next task. Multiple dependencies lead to slippage, so the P50 completion date for blasting Block F has slipped 1.5 days to halfway through January 26.

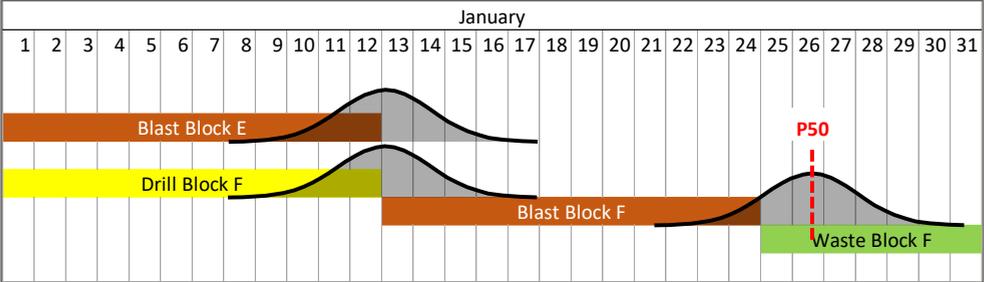


Figure 5 - 2 Blast Slippage

This slippage continues to grow throughout the schedule, highlighting that gains are lost, while losses accumulate throughout the system. Table 5 - 1 shows the growth in slippage throughout the schedule sequence.

Total slippage across the 72-day period of this schedule is 10 days, which equates to a 14% delay. There is also a 25% probability that the task won't finish until March 27, a slippage of 14 days, or 19%. However, what I found quite remarkable, was the plan output probability, the probability of railing on time as per the deterministic schedule date of March 13 **is only 2%, in other words a P02.**

So, in this example, P50 in definitely does not equal P50 out!

Activity	Slippage (Days)	Slippage (%)
Drilling	0	0
Blasting	1.5	6%
Waste	3.5	10%
Coal	6	13%
Processing	8	13%
Railing	10	14%

Table 5 - 1 Schedule Slippage

Alternatively, let's say that mine management understands there is a difference between input probability and output probability and they actually want a mine plan that achieves P50 for the plan outputs. To achieve that in the above example would mean that the plan inputs required would have to have a 72% probability of being achieved in order to compensate for the time slippage. So **P72 in, equals P50 out**. Now, let's understand what P72 inputs mean. It means that for a task that would take 10 days using mean input assumptions (deterministic), it is actually going to be scheduled to take 11.8 days (or 18% longer) to allow for the slippage that's going to occur. Now that sounds fine, until the production team and site management understand that you're scheduling tasks to take 18% longer than target. That won't sit well, as there will be a belief that the plan is driving the wrong behaviour in lowering productivity expectations.

The above analysis highlights why we build inventory between tasks in mining operations, because it increases the probability that equipment can continue to operate and is not held up by the preceding activity in the sequence. Let's run this same example but this time introduce inventories, allowing us to understand the impacts of inventory. Mine sites vary significantly in the inventories they hold and not all of that inventory is between tasks within the schedule. But, it is not uncommon for the blast crew to commence loading the shot as soon as it is finished

Activity	Slippage (Days)	Slippage (%)
Drilling	0	0
Blasting	0.5	2%
Waste	1	3%
Coal	2	4%
Processing	3	5%
Railing	4	6%

Table 5 - 2 Five Day Buffer Schedule Slippage

Next I found myself wondering, what is required for a P50 input plan to result in a P50 output plan? So I ran a number of variations of the above example and varied the time buffers, those results are shown in Table 5 - 3. The last row of this table is the input assumption probability required to achieve a 50% probability of achieving output assumptions. At buffer lengths of 13 days we still only have a P44 output plan and the minimal increase in P out between the 10 day and 13 day cases implies that buffers would have to be very, very large before P50 in equals P50 out in the plan.

Parameter	Unit	Length of Buffers					
		0 Days	3 Days	5 Days	8 Days	10 Days	13 Days
Schedule Time Slippage	Days	10	7	4	2	1.5	1
Schedule Time Slippage	%	14%	10%	6%	3%	2%	1%
Output Probability (P Out)	%	2%	17%	29%	38%	43%	44%
Required Input Probability (P In)	%	72%	62%	58%	55%	53%	52%

Table 5 - 3 Simulation Results

It is interesting to note that, with the introduction of inventories, only a very minor shift in the input assumption probabilities is required to achieve a P50 out plan. For example, with buffer lengths of 8 days, P55 input assumptions will result in a P50 output plan. With the parameters used in this example, P55 input means you're going to schedule the task at 3.5% below its expected (P50) production rate.

This example was built on activities that all had uniform task lengths of 12 days per task and with constant time buffers between every task. This is obviously an oversimplification of mine scheduling where task lengths vary substantially, as do time buffers between activities. However, at 13 days of time buffer, the buffers are actually greater in length than the actual task times (12 days) and we are still only at P44. So, I think it is safe to say that unless you are carrying massive inventories:

P50 in DOES NOT equal P50 out!

What really strikes me after doing the work to write this chapter is how much we under-rate slippage in the mining industry and the degree to which we underestimate its impact. When I think back about it, schedule slippage is not a term that I've ever heard used when discussing schedules, it is not something that I have ever seen actually allowed for in any schedules and it is not a KPI measure that I've ever seen in any mine plans. I suspect this is largely because we (mining) have always lived in a deterministic world and we don't have the scheduling software to easily identify, quantify and allow for slippage.

But what I really find incredulous is the lack of attention (that I've ever noticed anyway) to an issue that has such a large impact on throughput. If the mine is running with minimal inventories, then the ore railed is likely to be in the range of 5 - 15% less than planned. Let's take an alternate scenario where all of the mine site equipment was operating at 10% below expected productivity leading to 10% less ore than planned, do you think that would be noticed? Do you think that would get some management attention, with significant efforts to solve it and a recovery plan put in place? Absolutely!

However, we can have 10% slippage in our mine plans, plan after plan and it sits quietly under the radar. Why? I believe because we don't understand it, we don't

measure it, and with the high frequency of changes to plan and other inherent mine planning issues, we're not even aware that we are suffering 10% slippage in our mine plans. This is one of the significant benefits of probabilistic (or stochastic) scheduling, they are based on a deterministic schedule as the underlying plan. So any slippage is immediately evident, can be measured and we can start to bring this issue to light.

What Is The Solution To Schedule Slippage?

Famously, once we humans have learned and conquered something new, we can't go back to what we used to do. A perfect example of this is small children who can crawl, once they learn to walk and run, they can't go back, they don't go back, to crawling.

I've found it's the same in the mine scheduling world. Once you start thinking about stochastic scheduling and working with it, you can't go back. You can't go back to deterministic scheduling. I personally have been amazed at the step change in my thinking since I've started on this journey of advocating for change away from deterministic thinking towards stochastic thinking.

For many years now I have thought that we should be stochastic scheduling. But as the scheduling software out there hasn't had that capability, it has been left in the realm of simulation software instead and I've never spent too long dwelling on it. Recently, my thinking has changed as I have continued to question why we schedule deterministically instead of stochastically. I think that deterministic scheduling is one of the biggest downfalls in mine planning today, particularly given our coding and hardware abilities. I really do fail to see why stochastic scheduling capability is not industry standard in our mainstream mine scheduling software.

And as I said, once you start thinking about it, you can't go back.

This experience has been an awesome transformation for me, even at my tender age of 56 with over 30 years in the mining industry. Following the research carried out and examples I have compiled for this book, I now appreciate even further the pure folly of deterministic scheduling. **It really is the root of all evil in mine planning.** We have such a complex set of processes in mining, with so many activities, so

many dependencies between them and an incredible amount of uncertainty in every task. And we think that we can create a schedule that is a single snapshot in time, with one clearly defined plan and outcome - that is simply **crazy**.

What is even more ridiculous is that we put so much faith in these deterministic plans that we measure the performance of them and we measure people against them. As discussed earlier in this book, compliance to plan, that is measuring if equipment is in the right location at the right time, is a terrible measure. It places a blind faith and belief in our mine plans that is undeserved and shows a complete lack of understanding of the variability that exists in mining processes.

So, having ceased to agree that deterministic schedules have the high value we place on them, I now find my thinking has moved into a totally different sphere and I believe a significantly more valuable one at that. Here are some examples of the questions I now think about, some of which I have reflected on in past LinkedIn articles and some of which will be discussed in future LinkedIn articles:

- Where in the process should we carry inventory?
- Which is more cost effective for reducing slippage, carrying inventories or reducing variability?
- What is the real cost of not meeting plan?
- Should we be using the critical path approach in mine scheduling?
- Slippage is definitely a critical issue in short term plans, but how critical is it in medium term and long term plans?
- Is it flexibility in the mine schedule that allows us to get away with deterministic scheduling?
- What is the cost of having this flexibility in our planning?

I note that these examples are just the tip of the iceberg. Now that I've "learned to walk" and started thinking about these questions, I'll only continue to think this way and will find many other issues to explore. I firmly believe we need more people thinking this way in the industry and challenging paradigms, it is for the good of the industry. I hope this book has either started this journey for you, or spurred you to up the ante on an industry move to stochastic scheduling.

Now you know mining tasks have a high degree of variability around how long they take and want to build that into your schedules, what should you do? Well, Liam (a

reader of one of my LinkedIn articles) asked me a couple of questions, the first of them was, “Is there any software that makes weekly scheduling more stochastic? Or is just being realistic the crux of it?”

Great question.

Unfortunately, I haven't come across any software that includes stochastic functionality, which in my opinion is ridiculous. So what do we do about it while we're waiting for that capability?

Liam is right on to it because stochastic scheduling is not only about producing a schedule in software with ranging around the task times, but it is just as much about the mindset. And not just your mindset, but the mindset of a large range of people at your mine. It is about stepping over the hurdle of assuming that a deterministic plan is right. It is about ceasing to think and talk like you can actually hit those milestones that are dictated in the deterministic schedule. It is about understanding that there is variability inherent in the execution of a schedule and you'll never achieve 100% execution of the plan.

Therefore, it is not a matter of fixed time slots for each task, but, instead, realising that every single task within the schedule will have a range of times within which it will both start and finish. This will then impact the interaction between tasks, and, therefore, whether the timing of downstream tasks is further affected. While a schedule with every task scheduled on a stochastic basis would highlight that beautifully, we don't currently have that capability.

So what is it that we can do to help shift the mindset?

I believe the solution lies in manually implementing what the schedule should automatically be doing. My recommendation is to choose the most critical tasks within the schedule, then carry out your own ranging around the possible start and completion times for those tasks. The person who carried out the schedule will inherently have a good feel for which tasks are the most critical. The items that I would consider when determining those critical tasks are:

- First, those tasks with the lowest lag between tasks and, therefore, with the highest probability of the second task being delayed by the first task; and,
- Second, those tasks that are critical in the flow of the schedule achieving its key targets. For example, often in a schedule there are key blocks of ore

that are required for the blend, or a task that, if it runs late, is going to hold up multiple other tasks within the schedule.

With these critical tasks selected, then apply a range of finish dates to them. That range can be determined by analysing historical production rates for the relevant equipment and subsequently determining the potential variability. That variability will allow you to work out a range of times within which that task might finish. Then, instead of talking specific dates that the task will finish (like we do with deterministic schedules), start referring to the range of times within which it might finish. Better still, also start referring to the probability that it will not finish by a constraining date, creating an issue with a task downstream of this interaction, for example, there is a 40% probability that we won't finish drilling that block before we need to start loading the blast. The key element here is the language that will be used, it is not a language about specific end dates, but is one of date ranges and percentage likelihoods of there being an issue.

Mine plans are not just about working out how the mine will achieve certain targets and the sequencing of equipment required. A plan is also about managing risk and risk is caused by uncertainty (variability). Talking in a language of probabilities that an event will occur is identifying the risk and putting it squarely on the table for monitoring and management.

I recommend choosing a fixed number of the most critical items and carrying out this process every short term schedule. How many critical items are selected depends on how much time you have available, as this step does increase the labour time required to produce each schedule. You might have a Top 3, a Top 5 or even a Top 10 of critical tasks. But choose a fixed number and start to manually carry out this process and start the transition to talking in a different language.

I could be wrong, and I might need to further test this hypothesis, but at this time I don't believe this manual process is necessary for mid term or long term schedules.

Manually carrying out the ranging process will also assist the transition to schedule automation when that time comes. Why? I've always been a big believer that to get process automation right, you need to have the depth of understanding that comes from having done it manually yourself beforehand.

That's my recommended solution until such time as suppliers can provide us scheduling software with stochastic capabilities. Let me know how you go.

6

CRIME 3: Great Mine Plan – What Do We Do With It Now?

The Grade Control Superintendent role at Paraburdoo provided me with awesome learning opportunities. As discussed in the opening story to this book, we built an iron ore stockpile every week that had to achieve eight quality targets, this stockpile was called a “run”. Paraburdoo was a 24-hour, 7-day per week operation with a four-panel roster, so there were four grade controllers on my team, one on each shift. Shift lengths were eight hours, so at any point in time, three grade controllers were rostered on shift and one controller was on days off, which meant that throughout the length of the 7-day run construction, there would be four different people accountable for the implementation of the plan. We weren’t just responsible for the implementation of that plan, we also built the plan in the first place and modified the plan as required. As we built the run, we would get continuous feedback on the qualities of the ore that we had fed to the milling plant to date.

So, we had four separate people accountable for both creating a plan, modifying the plan, and implementing the plan. Over time, I found that this didn't work. Renowned business coach Verne Harnish has a saying, “if there is more than one person accountable, then no one is accountable,” and this proved to be true. We had four people trying to manage the same plan and all being held accountable for it. But, four people accountable for the plan always worked out to be no one accountable for the plan, because they could just blame one of the other grade controllers.

To solve this, I implemented what I called the “Run Master” system. Before each run started, I selected the Run Master for that particular run and notified all relevant parties. The selection process was typically the grade controller that was on day shift for the majority of that run construction, as they had the greatest opportunity to interact with other key personnel who had a role in mine plans, such as geologists, surveyors, maintenance planners and process engineers. The Run Master was the nominated person who was accountable for that weekly plan.

The Run Master still needed the other grade controllers to help them and those other grade controllers would often reschedule throughout their afternoon or night shifts. However, it was always done in consultation with the Run Master and no change to the plan was ever implemented without their approval. This system worked brilliantly; it was one of the best system changes I've ever implemented in terms of clear accountability. For each run, the clear boundaries on supervision, monitoring and management led to significant performance improvements in implementing the plans and achieving better qualities.

What Is The Problem?

Have you ever sat in one of those weekly planning meetings that drag on and on and on, with each engineer getting up and talking about their pit or their process, such as drill and blast, they talk through every activity, one engineer after another? The meeting goes on for an hour or longer and there are people falling asleep. It's like pulling teeth and no one walks out pumped up and ready to make it happen. Worse still, at the end of the meeting, no one is wiser as to the real issues that they need to focus on. Production personnel know the equipment sequence, but they don't know the importance of that sequence and where the mine plan might fall over.

Having worked at over twenty mine sites, I have had the unfortunate experience of attending numerous planning meetings just like this. So why does this happen? Two primary reasons spring to mind:

1. Mine planners put a lot of work into their schedules and they are passionate about them. So, it's only natural that they want to tell everyone how complex and technical the schedule is, in every little detail, and what a fine job they have done to actually manage to get a schedule to work. After all, they are engineers and engineers thrive on analysis, logic and details.
2. Mine planners think that those charged with implementing the weekly (or daily) plan need to know all this level of detail so they can do their job. This is their “handover” process from the planner to the executer.

So, once the schedule is complete, has been communicated via the weekly planning meeting and sent out to a cast of thousands via email, it should be easy for everyone to just “follow the plan” right? **Absolutely not!**

Mining operations are loaded with moments in time where decisions are required. Some of those decisions are not so important, but some are critical. The question is, “have you given the decision makers the right information, and enough of it, to allow them to make the right decisions?” This is particularly the case when you're not available to assist, such as on the weekend or in the middle of the night. In my experience, rarely have the executers been adequately armed with the right information.

The most common form of communication of a mine schedule is a Gantt Chart, an example of which is shown in Figure 6 - 1. A Gantt Chart might show the sequence

for all equipment, but that doesn't provide sufficient information when required to help those charged with executing the plan. You haven't given them enough context or background on the critical issues so that they are armed to make the right decisions.

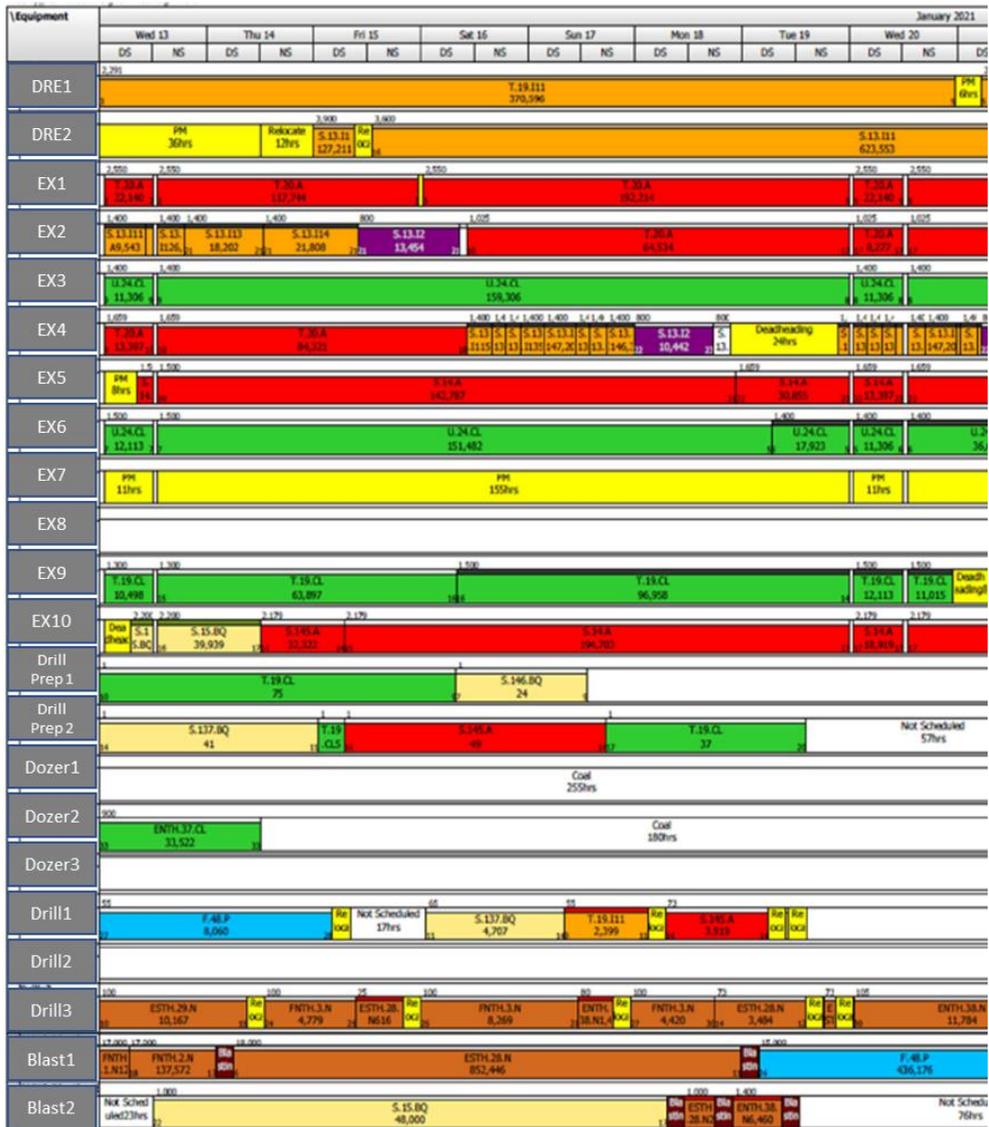


Figure 6 - 1 Example Gantt Chart

Worse still, there is an expectation that your mine schedule, which is but a single snapshot in time, can be executed as you have planned. If you truly believe this, if you really think that a mine schedule as complex as the Gantt above can be implemented, when it is built on a set of singular assumptions for so many parameters in the mine, then, can I please ask that you read a book called *The Flaw of Averages* by Sam Savage.

Plan the Work, Work the Plan.

I've heard this phrase or something similar many times. Mine sites follow this mantra because they don't understand variability and there is a belief that the schedule accurately reflects what will happen at the mine site. Consequently, one of the KPIs that is commonly in place within the mining industry is a measure called "compliance to plan". When it comes to compliance to plan, there are generally two types of compliance measured, the first is spatial compliance and the second is quantity compliance.

Spatial compliance is an attempt to check whether you mined from where the schedule said you should mine, so what percentage of the actual activity was in the plan. On face value, spatial compliance to plan seems like a good measure to have in place, but it can be very time consuming for surveyors and engineers to produce the results. I believe that monitoring spatial compliance shows a general naivety towards mine scheduling issues, the likes of which are discussed in this book. I would argue that this measure involves a lot of wasted labour that could be better invested elsewhere.

Spatial compliance to plan can be misleading if one activity has a high compliance to plan and the preceding or following activity has a low compliance to plan. That's an issue and it's going to get you into trouble. If something changes and one activity can no longer follow the schedule timing for that activity, then you need to review the whole schedule in light of this issue. There's a high likelihood that other activities should be modified as well, so that they no longer follow the original schedule.

Volumetric compliance is not concerned with where the activity was carried out, just whether you achieved the planned quantity. As further discussed in Chapter 7, I believe this is a poor measure to have in place and drives the wrong behaviour.

Why Is Poor Execution Of The Plan A Problem?

The best plans in the world are useless if you don't know how to execute them, because they aren't executed well. Poor understanding and execution of a mine plan has the following negative impacts on mining operations:

- Inefficient mining processes;
- Low confidence in the mine plan;
- Lost shipments or demurrage costs; and,
- Damaged customer reputation.

Let's unpack each of these impacts a bit more, starting with inefficient mining processes. A poor mine schedule or poor execution of it leads to changes to the plan during execution. This results in inefficiencies that include:

- Idle time due to one activity having to wait for another activity to finish.
- Idle time due to one activity interfering with another, such as the electric cable running across the intended haul path.
- Time lost deadheading as an alternative to idle time due to one of the above issues.
- Taking shortcuts that result in lower productivity or higher costs, just to keep operations continuing. Examples include:
 - Firing only part of the planned blast so that equipment has material to dig.
 - Mining coal to an intermediate stockpile due to the need to get it out of the pit.
 - Excavators digging half width or half depth faces.
 - Draglines switching early from extended keys to blocks to uncover coal.

Every one of the examples above presents a lost opportunity, you will not get that delay time back to be operating or to excavate at full productivity instead of reduced productivity. Every time a mine operates inefficiently, that costs the mine site money and those losses over the period of a year can sum to a very large total cost.

Low confidence in the mine plan is a disease that many mines are infested with. It becomes demoralising for mine planners to produce schedules that they know are a waste of time as they're not going to work. Teamwork becomes very challenging with operational personnel who think that your work is a waste of time. And for operational personnel, they are charged with implementing a plan that they have little faith in, which commonly extends to having little faith in everything else provided by the mining technical team, such as designs. Low confidence in the mine plan usually contributes strongly to a poor culture at the mining company.

The end goal of a mine schedule, and what determines whether the schedule is acceptable or not, is whether sufficient ore is produced to meet the shipping requirements. There is typically a long sequence of activities leading up to railing and shipping, so there is an increased probability that poorly implementing the mine schedule will result in either insufficient ore being produced or the wrong quality ore. The resultant shipping demurrage costs can add up to a large sum very quickly.

There are two key criteria that mine customers consider very important from their suppliers, these are:

1. Guarantee of supply; and,
2. Quality of the product.

Failure to meet a shipment, or sending an off-specification product to a customer, harms the reputation of the supplier. It weakens the supplier's bargaining power and could come at a very high cost.

Let's be blunt here, execution of the plan is a big part of where the money is won and lost at a mine site.

The Solution To Executing The Plan

You might have fantastic mine planning in place, but, if it's not being communicated well, even the best mine schedule in the world is rendered useless. So, the solution to this issue begins with having a full understanding of what a mine schedule is, why mine schedules are carried out and what is important in the schedule and what is just noise. Key to executing the mine plan, is understanding the variability that exists

within schedules (as discussed in Chapters 4 and 5) and that it is never going to be possible to execute the schedule exactly as it is on paper.

Mine schedules provide guidance as to the sequence of activities that need to be carried out and an indication of how long each activity should take. However, that is not where the real value lies in a mine schedule. The value is not in what happens during those activities but in understanding what happens at the boundaries of those activities. The nuggets of gold come in understanding and managing the schedule interactions. By schedule interaction, I mean where there are dependencies between activities, such as:

- One activity needs to be completed before another activity can be started on that same block, for example blast the block before it can be excavated.
- Two activities cannot operate in close proximity to each other on different mining blocks, for example the drill can't work on a bench immediately above where the dragline is working on a lower bench.
- There is another issue causing one block activity to be dependent on another block activity finishing first, for example drilling one block will stop material from another block being hauled through that area.

The critical issues surround the details of the interactions, i.e. what is the preceding activity that needs to be finished prior to the interaction, what is the activity following the interaction, and, most importantly, what is the lag or float between those activities. It is these interactions that we need to identify, quantify and manage. The real trick to executing a schedule well is managing these interactions.

Think about it, if every interaction was managed so that they stayed within the target window, then activities would never overlap and there would be no need to constantly re-run the mine schedule. Equipment wouldn't suffer from unplanned idle time or deadheading, there would be no need to take inefficient shortcuts and, subsequently, the overall productivity of the mine will improve dramatically. However, just because you manage the interactions and no longer have overlaps, that doesn't mean that you will automatically achieve all the schedule targets. For example, if your schedule assumptions are uniformly too optimistic for all activities, there might not be any schedule conflicts, but you are unlikely to achieve the ore targets of the schedule.

Here is how I believe a mine schedule should be executed. It doesn't need any fancy planning software; it can be done with a spreadsheet. However, the right software in place would help the mine planners as it will speed up the process, giving the engineers time back to invest in better scheduling and smarter mining.

The other issue to point out is that all of the following discussion about mine plan execution only relates to your shortest term mine schedule. At most mine sites, this will be a weekly schedule, but it could be a daily schedule or similar at various mines. This process only applies to the shortest-term schedule because that is the only schedule that is actually executed, all other schedules are purely for information or decision-making purposes.

Once your schedule is complete and ready to communicate, here is what you do.

Start with selecting the most critical interactions, the ones that will have the greatest impact on the mining operations if there is a delay. I will explain more on how to do this later. Start simple. I recommend choosing the Top 5, then, for each of those interactions, focus on the two activities either side of the interaction.

To help explain the process, I am going to use an easy example. Let's talk about an excavator digging Block B, which it will start after it has finished excavating Block A. However, before the excavator can start on Block B, that block must first be blasted by the blast crew. So, blasting is the activity preceding the interaction and excavation is the activity following the interaction. Schedule management is the term that I am going to use within this chapter to describe the process of executing the schedule. Schedule management involves ensuring that an overlap does not occur at this interaction, so no equipment needs to be parked up or relocated.

The activity preceding the interaction is blasting so schedule management starts with understanding and monitoring what needs to happen with that blasting activity, and ensuring it is not sufficiently delayed as to create an overlap with the excavation. If the lag between the two activities is three days, then the critical number to determine is the minimum rate that the blasting activity must be carried out at so that it is not going to take three days longer than planned. A slower rate than this will use up the float and result in a delay to the excavation.

As shown in Table 6 - 1, this involves considering the activity quantity and time available, which leads to the minimum rate required per shift. Note that the

minimum rate can be over any time period that is suitable, so it could be per day, for example. But I think that per shift is most relevant, so each Supervisor and their crew know what is required.

You can see in Table 6 - 1 that if the blast crew load at less than 45 tons per shift, then the activity will be sufficiently delayed so it won't be blasted in time and the excavation process will be held up. In the last row of the table we are also continuing to monitor the expected lag, so we can understand if the situation is improving or deteriorating over time.

Parameter	Unit	Value
Current Date		11-Feb
Required Date		19-Feb
Shifts available (Day Shift only)		8
Explosives Remaining to Be Loaded	t	360
Minimum Load Rate Required	t/shift	45
Expected Load Rate	t/shift	70
Expected Finish Date		16-Feb
Expected Lag	days	3.2

Table 6 - 1 Preceding Activity Requirements

The preceding activity is only half the interaction, you also have to monitor the activity following the interaction. In this case, the following activity is the excavation, with the excavator digging Block A before commencing excavation of Block B, which is the block that has just been blasted. If the excavation of Block A is ahead of plan, then we may have an issue as the probability of an overlap of the two activities has increased. Alternatively, if the excavation of Block A is behind plan, then the probability of an overlap has decreased.

Again, the key number is the rate that would create an overlap, although this time it is reversed and we are determining the maximum rate rather than the minimum rate. So, if excavation is carried out at higher than this rate, then the lag between activities is reducing and an overlap may be looming.

Of course, this is not to say that the excavation shouldn't work at a rate higher than the maximum target, as high productivity is key to maximizing profits. It is about the awareness that if excavation does work at a higher rate, then either the excavator will park up or have to relocate to another block, or the blasting has to be sped up somehow.

As shown in Table 6 - 2, if excavation achieves more than 19,100 bank cubic meters per shift, then the float reduces and the probability of an overlap increases.

Parameter	Unit	Value
Current Date		11-Feb
Expected Available Date		16-Feb
Shifts available		10
Remaining Excavation Volume	bcm	191,428
Maximum Excavation Rate (No Overlap)	bcm/shift	19,143
Expected Dig Rate	bcm/shift	11,500
Expected Finish Date		19-Feb
Expected Lag	days	3.2

Table 6 - 2 Following Activity Requirements

Schedule management involves monitoring this interaction and recalculating the updated numbers on a regular basis. I suggest monitoring at least daily, but it would work best if it is done on a shift by shift basis. Table 6 - 3 shows a series of tables

highlighting how this process might work over a period of three consecutive shifts. Notice that, as the blast loading is underperforming and the excavation is overperforming, the lag shrunk from 3.2 days to 1.8 days over the three shifts.

Block Number	B81-35
Date	11-Feb
Shift Completed	Night

Parameter	Preceding		Following	
	Unit	Value	Unit	Value
Activity		Blast		Excavation
Quantity Completed Previous Shift	t	0	bcm	14,337
Quantity Remaining to Be Completed	t	360	bcm	191,428
Required Date		19-Feb		16-Feb
Shifts Available	D/S Only	8		10
Minimum/Maximum Rate Required	t/shift	45	bcm/shift	19,143
Expected Production Rate	t/shift	70	bcm/shift	11,500
Expected Finish Date		16-Feb		19-Feb
Expected Lag	days	3.2	days	3.2

Block Number	B81-35
Date	11-Feb
Shift Completed	Day

Parameter	Preceding		Following	
	Unit	Value	Unit	Value
Activity		Blast		Excavation
Quantity Completed Previous Shift	t	16	bcm	18,965
Quantity Remaining to Be Completed	t	344	bcm	172,463
Required Date		19-Feb		16-Feb
Shifts Available	D/S Only	7		9
Minimum/Maximum Rate Required	t/shift	49	bcm/shift	19,163
Expected Production Rate	t/shift	70	bcm/shift	11,500
Expected Finish Date		16-Feb		18-Feb
Expected Lag	days	2.6	days	2.6

Block Number	B81-35
Date	12-Feb
Shift Completed	Night

Parameter	Preceding		Following	
	Unit	Value	Unit	Value
Activity		Blast		Excavation
Quantity Completed Previous Shift	t	0	bcm	17,498
Quantity Remaining to Be Completed	t	344	bcm	154,965
Required Date		18-Feb		16-Feb
Shifts Available	D/S Only	7		8
Minimum/Maximum Rate Required	t/shift	49	bcm/shift	19,371
Expected Production Rate	t/shift	70	bcm/shift	11,500
Expected Finish Date		16-Feb		18-Feb
Expected Lag	days	1.8	days	1.8

Table 6 - 3 Schedule Management Example

If you carry out the schedule management process for the Top 5 interactions, you are not only going to be focusing on the right areas to allow you to implement the schedule, but you also now have something to build your planning communication around. Let me ask this question, if you were starting a brand new mine today and were incorporating the Top 5 process, how would you build your planning system? I think it would look very different to what is happening at your mine site now.

Here is how I would design the mine planning system.

First, the Top 5 are going to be a driver when I am scheduling. I'll be aware of the overlap probabilities during the scheduling process and looking to minimise them. As I want to make the schedule as easy as possible to execute, I want to build a schedule that is low risk. My scheduling process is going to be set up so that it is reporting the schedule interaction overlaps during generation of the schedule and producing what I would call a robust plan.

Once I've finished the mine schedule, now the weekly schedule meeting changes completely. It is no longer the boring monotony of talking through every activity with people falling asleep. Instead, it becomes a meeting of just the key players at Superintendent level with the following agenda:

1. Identify the critical interactions and agree on which are the Top 5.
2. Discuss each of the Top 5 individually, how realistic the schedule is and whether we think the actual situation might play out better or worse than scheduled.
3. Discuss whether there are any roadblocks or critical issues we need to be aware of and how we lessen their likelihood or impact. A common example here is shared resources such as dozers, now there is a decision-making tool for determining the priority for those resources.
4. Decide what we need to do to ensure the overlaps don't occur.
5. Discuss the contingency plans if things do go wrong, if an overlap does occur e.g. is the equipment relocated to another block or is there another, smarter solution?

The Top 5 is also a great tool for facilitating the following discussions:

- Between Managers and Superintendents of different functions when trying to determine their team priorities (e.g. maintenance priorities), or when there are conflicts.
- Between Superintendents and their Supervisors in shift planning and monitoring how the operation is going.
- The shift handover between Supervisors, there is a clear discussion required around progress on the Top 5 and potential issues for the incoming Supervisor.
- Your daily planning meeting is now sorted as well. It is a quick review of performance on the Top 5 and why any of them worked well or why they didn't. As well as a chance to address whether there are any looming issues and come to an agreement on the solutions. It's at this meeting that you start to identify issues early and discuss contingency plans if they are required.

With numerous schedule interactions in every schedule, how do you decide which of them make the Top 5? My suggestion is to consider three parameters:

1. The length of the scheduled lag;
2. The variability of the activities either side of the interaction; and,
3. The level of impact on the mine if the overlap of activities does happen and one activity is modified or delayed.

The first parameter of the length of scheduled lag is easy, it comes straight from the mine schedule just completed.

The second parameter to determine is the variability of activities either side. It is rare for the daily production to be normally distributed (as discussed in Chapter 4), but we'll use that assumption in this example as a simple proxy for variability of the activities. Using historical productivity data, calculate the standard deviation as a percentage of the daily production for the activities each side of the interaction. Then, average the two standard deviations and use the whole number version of this average in the calculation.

Finally, for the level of impact on the mine determine a score between 1 and 10 that represents the level of impact if a schedule conflict does occur. Where 10 is an

extremely high impact, such as missing a shipment, and 1 is extremely low, such as having to relocate an excavator to another dig location close by, i.e. it will have minimal impact on ultimate production levels.

To determine the Top 5, divide the standard deviation (parameter 2) by the lag (parameter 1) then multiply by the impact (parameter 3) and you'll have a way to compare all the interactions within the mine schedule. The higher the number, the bigger the issue and the more likely it should be in your Top 5. An example calculation is shown in Table 6 - 4, the number itself is not important, it is the relativity between the numbers for each interaction that is important. This is a very simplistic calculation process, but it is indicative of which interactions should be considered for inclusion in the Top 5.

Parameter	Unit	Value
Preceding Activity Mean Daily Production	t	61
Preceding Activity Standard Deviation	t	23
Standard Deviation (% of Mean)	%	38%
Following Activity Mean Daily Production	bcm	12,365
Following Activity Standard Deviation	bcm	5,974
Standard Deviation (% of Mean)	%	48%
SD Average Expressed as Whole Number		43
Schedule Lag	days	4
Impact		7
Score		75

Table 6 - 4 Top 5 Indicator Example

If you get the Top 5 process up and running smoothly, you could always try expanding it to the Top 10. However, I would caution against it, I have always been a big fan of the maxims “keep it simple” and “less is more”.

7

CRIME 4: You Have KPIs, But A Lack Of Mine Planning KPIs

One of the worst examples of a key performance indicator (KPI) design I ever came across was while employed at a mine site. KPIs are meant to drive behaviour, and they certainly can be very powerful at that, but unfortunately that may not always be for the overall benefit of the mine site.

The maintenance team had a KPI that was fairly common throughout the industry at the time, that is “mean time to repair”, and the team were measured and incentivised financially against this measure. Like a lot of KPIs, the measure had exceptions built into it for elements that were not under the maintenance team’s control. For example, if a piece of equipment was down for maintenance because of accident damage, that was not the maintenance team’s fault, so this downtime was not included in their mean time to repair calculation (i.e. they weren’t held accountable for the time taken to repair accident damage).

We had the scenario where a shovel that was down for maintenance, was down because of accident damage. At the same time, there was other equipment down for maintenance that was down because of mechanical or electrical issues. During an execution plan we desperately needed the ore that was located in front of the shovel that was down for accident damage. The shovel was parked right in front of the face so that ore could not be mined by any other excavation equipment.

However, because this particular shovel was down for accident damage, it did not count against the maintenance department KPI so it was not a priority for them. However, it was a priority for us! But as much as we pushed, we could not get that shovel prioritised for repair. I had to take this issue upwards through two levels of management before I could get the maintenance priorities successfully changed.

What Is The Problem?

Like many industries, the mining industry loves KPIs, there are a plethora of them implemented at mine sites globally and they exist across a wide range of functions. KPIs are prevalent across the areas of:

- Safety;
- Mine Production;
- Maintenance;

- Financial; and,
- Human Relations.

Here are just a few examples of common KPIs:

- Safety –
 - Total Recordable Injury Frequency Rate
 - Number of Near Misses
- Production –
 - Equipment utilisation percentage
 - Drill metres per shift
- Maintenance –
 - Availability percentage
 - Mean time to repair
- Human Resources –
 - Labour turnover
 - Days to onboard
- Financial –
 - Gross Margin
 - Cost of inventory

If a mine has KPIs for so many other functions, why then, when it comes to mine scheduling, is it complete silence out there? What mine planning KPIs do you know of? What KPIs do you have installed or in place at your mine site?

Compliance to plan is the only mine planning KPI I have regularly come across in the industry, with there being two versions of this, spatial compliance to plan and quantity compliance to plan. However, compliance to plan is a bad measure, it drives the notion that a deterministic plan will work and is actually capable of being implemented in accordance with the plan. As an industry, we have to give up on this notion of being able to exactly implement a mine plan as it has been calculated. Compliance to plan is not a measure of how good the plan is, nor is it a measure of whether the plan is high quality, it is actually a measure of how well the plan was executed.

If you are producing quality mine plans that have variability built into the activity completion times (as discussed in Chapter 4), then if you are measuring compliance

to plan and reporting compliance in percentages, what should you measure and report against? For example, is the mean the average that you should report against? Or is it the median or the mode? Figure 7 - 1 shows the level of difference that can occur between these numbers, so are any of them true indicators of compliance?

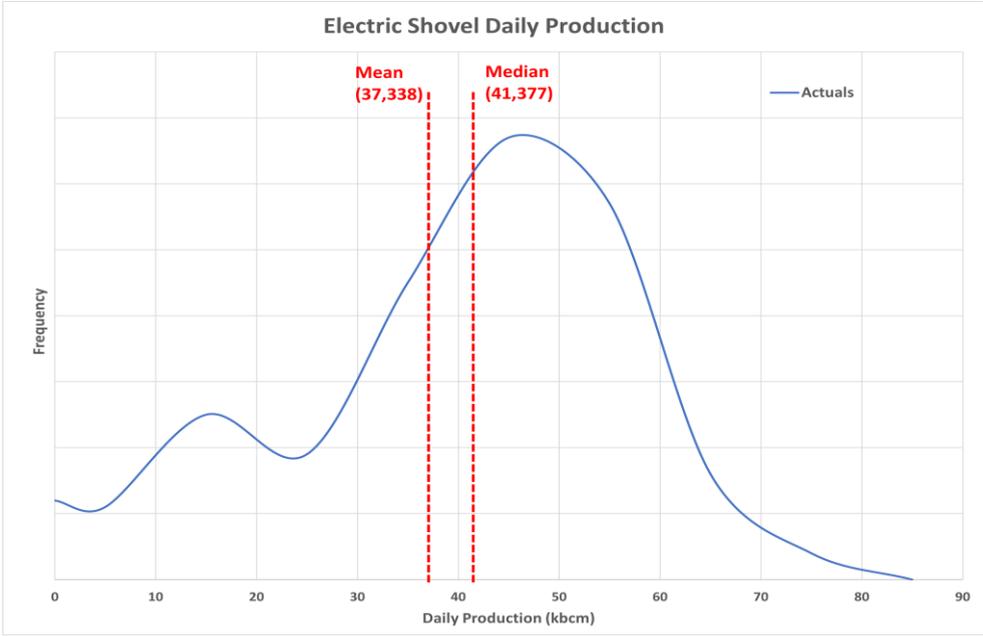


Figure 7 - 1 Median or Mean?

In addition, which person are you measuring the performance of when you measure compliance to plan? Are you measuring the performance of the mine planner? Or the personnel who are charged with executing the plan? If you have this KPI in place, who are you holding accountable for it?

Why Are Lack Of Mine Planning KPIs A Problem?

Why is measuring the quality of your mine plan important? Think about these questions:

- Is your mine planning getting better or worse over time?
- How does your mine planning compare against other mines or your competitors?
- What is the impact on mine planning of swapping people between planning roles?
- Your current mine scheduler talks the talk, but do they walk the walk? Are they actually good at producing mine plans?
- If your mine scheduler wants to improve themselves and their planning, do they know how to? And can they measure when they have actually improved?

It would be very useful to know the answers to these questions. There is a debate that has been raging for as long as I have worked in the mining industry, particularly when things are not going well. It centres around the question, “was it a poor mine plan or was it poorly executed?” Without being able to measure whether or not you have a good mine plan, that is without having the right KPI in place for your mine plans, it is not possible to answer that debate as you don’t know whether it was a good mine plan.

It may help if we sit back and think about why we have KPIs in place and how they work. KPI’s allow us to pick some key measures to focus on and to set appropriate targets for those measures. Targets generally are derived via assessment of previous site performance or industry benchmarking. Progress is then tracked and reported against that benchmarking. Trending is an underrated, but very important part of KPI systems, and by “trending” I mean tracking the change in performance over time. Measuring and monitoring trending is a critical component in continuous improvement.

Mine planning is what the mine site bases the whole mining operation around. So why don’t mine sites have KPIs in place that drive better planning and also allow

them to track trends and plan performance over time? How do you know if your mine planning is getting better or worse?

What Is The Solution?

As previously discussed, I don't believe compliance to plan is a good measure. Mine sites typically do have measures in place related to mine plans, but, from my perspective, plan related measures can be categorized into three types.

1. Those that measure the quality of the plan, which is what this Chapter is about.
2. Those that measure the quality of the execution of the plan, which is what I believe the compliance to plan measure is.
3. Those that measure the efficiency of production. These are outside the scope of this book.

What measures should you put in place to measure the quality of the mine plan? Here are two key measures that I would consider putting in place to track how mine planning is going and, in particular, how it is trending over time.

1. Plan Score.
2. Plan on Plan Change.

Every mine is different so it is not possible for me to detail what these measures would look like for your mine site in a generic book about mine planning exactly. Therefore, I will discuss them on a general basis.

Plan Score

It's time for another story.

At Paraburdoo, while building each stockpile, we had eight qualities we targeted in each stockpile construction, that stockpile construction being labelled as a "run". Those qualities were as follows:

- Lump – iron, alumina, phos and silica content

- Fines – iron, alumina, phos and silica content.

For each run, we were provided specific numerical targets for each of the eight qualities, based on the grades of ore already at the port and upcoming planned shipments to customers. While we were provided single point numerical targets, we weren't provided acceptable ranges around that target for each of those eight qualities. However, following many years of run construction, the Grade Control team had fashioned our own version of what we believed were acceptable ranges, for example:

- Lump iron could vary 0.2% either side of a target of about 62%.
- Fines iron could vary 0.4% either side of a target of about 58%.
- Fines alumina could vary 0.03% either side of a target of about 2.8%.
- Fines phos could vary 0.002% either side of a target of about 0.083%.
- Lump silica could vary 0.5% either side of a target of about 6%.

(note that this story is from 25 years ago and so the numbers are not factually correct and have been created for this example)

The first thing you may notice is the variation in how sensitive each of these quality parameters are. For example, the allowable variation on lump silica of 0.5% is 8.3% of the target of 6% - so the lump silica has an acceptable tolerance of +/- 8.3% around the target. However, the 0.2% allowable variation on lump iron is only 0.32% of the target of 62% - so the lump iron has an acceptable tolerance of only +/- 0.32% around the target. Or we could express this another way, lump iron is 26 times more sensitive to variation than lump silica.

One problem with this is that the Grade Controllers who scheduled the runs knew that some qualities were more important than others. It was not possible for them to focus on eight quality parameters at once while they were executing their plan. So, they typically focussed on three of them intently throughout the run construction, three of them they tended to keep a cursory watch on and there were two parameters that were monitored infrequently throughout the run.

This became patently obvious when we had one run that was an absolute disaster. The key qualities were pretty well spot on as the grade controller was heavily focussed on those, but the grade controller totally lost sight of silica and the lump and fines silica both finished up about 2% over their respective targets. The silica

was so far off target that the stockpile could not be blended at the port. I cannot adequately express in words how much grief this caused and the ramifications it had. The Port Logistics Team had to change the run targets mid run construction for the four other mines that we blended with, as well as delay multiple shipments. You can imagine how popular the Paraburdoo Grade Control Team were.

In response to this event, to ensure that all eight qualities were monitored, I created a new scheduling measure I called the Run Score. It was a score that, in one single number, represented how close the run was to being on target for all eight qualities. After many discussions with numerous stakeholders, I arrived at a weighting for each of the eight qualities. I then multiplied those weightings by the discrepancy between target and planned for each of the eight qualities, the sum of those eight numbers then provides the overall Run Score. Table 7 - 1 below shows an example set of calculations for the Run Score to help demonstrate the above explanation.

Schedule Parameter		Target	Plan Result	Discrepancy	Weighting	Score
Lump	Iron	62.0%	61.9%	0.10%	10,850	10.9
	Alumina	4.10%	4.11%	0.01%	2,392	0.2
	Phos	0.121%	0.116%	0.005%	706	0.0
	Silica	6.0%	6.1%	0.10%	420	0.4
Fines	Iron	58.0%	58.3%	0.30%	6,767	20.3
	Alumina	2.80%	2.82%	0.02%	4,900	1.0
	Phos	0.083%	0.084%	0.001%	1,453	0.0
	Silica	4.7%	5.2%	0.50%	411	2.1
Run Score						34.9

Table 7 - 1 Run Score Example

To arrive at the weighting factor in this example, I used the target for each quality divided by the allowable target band and then applied a constant modifying factor to all parameters so that a score of 100 was a fairly bad result. Of course, a score of

zero is a perfect result, so I then had a scale for my Run Score to try and operate within, i.e. all Run Scores should fall between 0 and 50.

I would note that this is just one methodology and you could build your own system with different logic. The Production Manager at the time wanted me to build a scoring system where 100 was the perfect score and 0 was a very bad result. I chose not to follow that path, because in implementing something extremely new such as this, I wanted it to be logical and easily transparent so it was widely accepted.

There were a number of unexpected bonuses that arose from this new Run Score mechanism. First, it became a very useful decision-making tool when actually creating the schedule. In our scenario we would normally create a number of different plans and then choose what we thought was the best of them. Before creation of the Run Score, it was difficult to determine which of the alternatives was the best option to implement. This was particularly so as it was impossible to achieve all eight qualities on target, so the Grade Controllers were frequently confronted with a decision as to which parameter should be off target and by how far.

The Run Score changed dynamics so that they didn't focus on any particular qualities, but simply generated options that produced a lower score. Even better, the Run Score provided some guidance as to where likely improvements could be made. In Table 7 - 1 it is evident that 31.2 of the score of 34.9 is driven by lump and fines iron discrepancy, so these are the two parameters to target to achieve a lower Run Score.

Second, the Run Score became a very useful tool during execution of the plan. Given the inherent variation in mine plans, frequent decisions are required during execution of a mine plan. For example, when it is no longer possible to follow the mine plan, should we take Path A or Path B? The Run Score provided an objective decision-making tool, the decision was simple, we chose the path that gave the better Run Score.

Lastly, it also provided a very effective KPI for tracking our scheduling performance. I could track whether our scheduling was improving each run, as I only had to look at the trend in Run Score results over time. Having also implemented the Run Master process where one of the four Grade Controllers on the 24-hour panel roster was in charge of each run (see start of Chapter 6), I could

then use the Run Score results for each Grade Controller to track how they were performing over time and determine who needed further coaching.

In the Paraburdoo Run Score example, I incorporated only our eight quality parameters in the Run Score calculation. I could have easily added in extra parameters that were important in the schedule, such as tonnages, but I didn't as I wanted to keep it fairly simple initially. However, a Plan Score could be created and introduced for any length of mine plan, from a daily plan through to a life of mine plan. It is simply a matter of selecting the important parameters in the mine plan and then using a weighting mechanism to arrive at a Plan Score. Some example parameters you might include are:

- Target ore tonnages.
- Equipment park up time due to mining interactions.
- Equipment deadhead time.
- How well a short term plan follows a longer term plan.
- Average lag between activities.
- Timing of capital expenditure.
- Net Present Value.

Now don't for a minute think this is a simple and quick process, it requires discussions with a large range of personnel as mine plans have many customers. Deciding the parameters to include and their respective weightings is a challenging exercise, as differing customers will focus on certain parameters and will place differing importance on those parameters. My recommendation is to start with fewer parameters initially and get the scoring system operational, then add more parameters over time if you think they're relevant. Keep it simple.

Plan Score may be a complex process to get right, but even if nothing else is gained, the learning from discussions with a range of mine plan customers to arrive at what they believe is important in a mine plan, and the relative importance of those parameters, is extremely beneficial. If you can achieve the end result of an operational Plan Score, you shouldn't underestimate how useful the mine planning tool you have just created is – for the areas of scheduling, schedule execution and monitoring and measuring.

Plan on Plan Change

Stability is very important in planning across any industry, it is hard to effectively create and execute a plan when the underlying activities have high degrees of variability. The best way to measure stability within mine scheduling is to reference the amount of change from one plan to the next. A high degree of change between two consecutive schedules suggests that the scheduling process is not in control and that there is something wrong in the process and, therefore, there is low stability within mine scheduling.

Consistent change between schedules leads to a lack of confidence in the schedule and, more importantly, a lack of confidence in the mine schedulers. It is hard to commit to implementing a schedule if you think there is going to be a lot of change between this schedule and the next revision of it. Constant change between mine schedules, from one schedule to the next, results in a lot of inefficiency within the mining process. Changes lead to additional time lost deadheading or operating with inefficient practices and it also means that it is harder to set the mine up for future tasks.

I believe a good schedule incorporates all of the 5 “Rs”, these stand for:

- Reliable
- Right (correct)
- Robust
- Rational (smart) and
- Realistic

One critical element required is robustness, a robust schedule stands up to the trials and tribulations of executing it. Or, as commented by one of my mining colleagues, “it survives first contact with the enemy”. Plan on plan change is a very good indicator of robustness, as a robust schedule will require minimal change from one schedule to the next, whereas high percentages of change between schedules implies that the scheduling process can be improved.

So how do you measure plan on plan change? Let’s use an example of a 6-week schedule that is rescheduled every week to help explain this process. For ease of explanation, let’s call the plan that was completed last week Plan 1 and let’s call the new plan that has just been completed this week, which we are going to measure the

plan on plan change KPI for, Plan 2. Because it is scheduled weekly, the period for the first week of Plan 2 is always going to be the same period as that of the second week from Plan 1. Similarly, week two of Plan 2 will line up with week three of Plan 1, week three of Plan 2 will line up with week four of Plan 1, etc.

It is then possible to calculate the percentage of material planned within week two of Plan 1, that is still planned in the corresponding period in Plan 2 (week one). This doesn't need to be a laborious process. It is a relatively easy calculation done by exporting the paths and quantities from both plans then carrying out a match between the two weeks in a spreadsheet. An example of this calculation process is shown in Table 7 - 2.

Block	Activity	Unit	Quantities Planned		Percentage Change
			Plan 1 Week 2	Plan 2 Week 1	
R44/S21/B06/D	Drilling	m	3,863	4,981	29%
R44/S21/B07/D	Drilling	m	2,754	733	73%
R39/S17/B01/C	Blasting	bcm	148,652	148,652	0%
R39/S17/B02/C	Blasting	bcm	117,437	117,437	0%
R39/S17/B03/C	Blasting	bcm	121,462	121,462	0%
R39/S17/B04/C	Blasting	bcm	112,692	34,632	69%
R18/S54/B11/C	Excavation	bcm	82,800	103,428	25%
R18/S54/B12/C	Excavation	bcm	112,450	101,537	10%
R18/S54/B13/C	Excavation	bcm	34,520	-	100%
R44/S21/B01/A	Excavation	bcm	17,344	17,344	0%
R44/S21/B02/A	Excavation	bcm	25,238	25,238	0%
R44/S21/B03/A	Excavation	bcm	23,174	17,355	25%
R44/S21/B04/A	Excavation	bcm	18,426	-	100%
R18/S54/B09/C	Excavation	bcm	-	3,867	100%
R18/S54/B10/C	Excavation	bcm	-	67,812	100%
Plan Average					42%

Table 7 - 2 Plan on Plan Change Calculation

This measure is a reflection of stability in mine planning at the mine site, the more stable the mine plan is, the more consistency from plan to plan and the lower the

level of plan on plan change. The impact this measure has on how well the mine site works cannot be overstated. Significant changes from one plan to the next are infectious and not in a positive way.

First, a high plan on plan change leads to a complete lack of confidence in mine planning across the mine site. If mine planning is of a high quality, then the plans last longer before they're wrong and need to be carried out again. This is a measure I like to describe as the "shelf life" of a mine plan. Lack of confidence in the mine plan is never a good thing, ultimately it leads to a desire on site to ignore the mine plan.

Second, high plan on plan change is an indicator of high inefficiency within the operations, the re-scheduling of the plan leads to extra walking of equipment and operation of the equipment in an unproductive manner.

Third, the operations team don't get the opportunity to set the mine up for productive phases and productive operations due to the constant change. The implementers of the plan can't do anything with confidence so they don't set themselves up for success.

It is important to be aware that the further out in the plan that is being analysed, the lower the percentage of plan change on change will be. The objective should be high consistency within the first two weeks, this is where quality planning is paramount and ideally there should be minimal change in the plan. The first two weeks are where execution of the plan is happening and it is where constant change will have the greatest negative impact. If your mine planning system leads to minimal plan on plan change, then you are a long way towards quality planning and execution.

8

CRIME 5: Bullshit In, Bullshit Out

Many years ago, “auto-schedulers” were in their infancy and the mine site I was working at had recently introduced a scheduling package used for mid term to long term scheduling. Changing mine design or scheduling software platforms at a mine site always makes for challenging times and exposes the mine planning team to the likelihood of errors, you just hope they’re not catastrophic.

We were in the process of producing our first five year plan with the new scheduling software and, of course, it was hectic. There are a lot of moving pieces in a mine plan and we were trying to pull them all together while trying to navigate our way through the intricacies of how the new software worked, which was slowing us down dramatically. Our lack of ability with the new software was putting us behind schedule for getting the plan out on time and that was piling the pressure on. We were working long hours, looking at the same numbers over and over and getting tired. Of course, that’s when things get missed.

We were using the auto-scheduler and, therefore, that required a series of rules around sequencing to be created within the software so the auto-scheduler created viable equipment sequences. In creating those rules for the first time, we missed creating a rule for when waste was dozed instead of excavated by a loading unit. This wasn’t a common task, but there was the occasional waste between coal seams that were too thin for dragline excavation so was dozed off the top of coal instead of excavated by dragline.

Turns out, in the second year of the plan, there was a strip in one of the pits that had an entire strip of dozer waste in it. Because there were no rules dictating the sequencing of that dozing, the auto-scheduler scheduled it right away at the beginning of the schedule. Why wouldn’t it, there was a lot of coal uncovered for very little work. So, in Year 1 of the plan, the schedule showed us uncovering 300,000t of coal that we weren’t actually uncovering until Year 2. We didn’t even notice the error when looking at the Gantt Chart, the waste quantity was small and the dozer task was just a small blip on a chart with periods of a year at a time.

In hindsight, it was funny, and you wonder how you could miss it, but when you’re under pressure and there are numerous issues to review and resolve, it is easy to miss one. It certainly wasn’t funny at the time. The error wasn’t picked up until after we had released the schedule and you can imagine how well it went down when we took 300,000t of coal out of the first year of the schedule.....

What Is The Problem?

Bullshit in, bullshit out. Have you ever heard that saying? I've heard it countless times in relation to a range of areas in life, but there is one area in which it is definitely applicable and that is mine planning. A plan is only as good as the data and assumptions it is built on. If they are wrong, the plan will be wrong.

There are two primary reasons that lead to planning inputs being wrong, they are:

1. An innocent mistake built into the plan.
2. A choice to use wrong data.

Let's tackle these one at a time, starting with innocent mistakes.

Mistakes happen, we're all human and nobody is perfect, so there will be mistakes that creep in to mine plans. The question is whether you find them before the plan goes to print. Whether you will find them or not depends on one thing – whether you carry out sufficient due diligence.

So why might this issue of a lack of due diligence occur? There are two reasons primarily, and both of them are relatively common:

1. The planner takes the schedule for granted.
2. The planner has insufficient time available to carry out the required level of due diligence.

Let's discuss the first of those, the planner taking the schedule for granted. This occurs because they get to know the plan too well and the whole thing becomes a bit “routine” for them, so they easily overlook things. On numerous occasions I have taken over someone else's schedule, particularly as a consultant for 13 years. I don't run a schedule without first getting to know the schedule and thoroughly checking that the basis is right. So, I have always started by carrying out a rigorous due diligence process before I start scheduling.

It is incredible the number of errors that I have found in schedules and some of them have been very simple ones at that. Most times I know the person who has been scheduling previously well enough to know that they are a very competent scheduler. They are a better scheduler than the errors I have found. I think in most instances those errors have come about because of over familiarity with the

schedule. At which time their brain sees what it expects to see, rather than what is actually there. Or the cause could have been the second issue – lack of time.

Let's talk about the other reason for poor due diligence, insufficient time. Mine sites roll out an incredible number of schedules on a very routine basis. It is very common that a range of actions are required to happen prior to the planner being able to commence, leaving insufficient time for the planner to carry out sufficient due diligence on their plan. This is particularly the case in short term and execution plans, where numerous data inputs are required from other site personnel. This issue can be further exacerbated where the planner is too nonchalant about the process because they've done it all before.

Mine plans always have a deadline, they have to be in circulation by a certain day and time and often this is driven by a regularly scheduled meeting at which the mine plan needs to be presented and discussed. The mine planner has numerous steps to take before they can even start scheduling, items such as updating face positions, updating the time usage model with latest maintenance information and inputting new design information such as block quantities. All of these steps, plus collecting inputs from others and an impending deadline for completing the schedule, leave the planner with insufficient time.

Why is it easy for the mine plan to have errors that go undetected? Let's just look at a typical mine plan and review how many places there are that errors could creep into the schedule and not be found. I have detailed below the data and input assumptions that go into a typical plan where errors can creep in.

The Underlying Reserves

Most mine sites break their pits down into blocks of some form, which is a division of the pit on a plan view basis, and then each block is further divided into horizons or benches. Alternately, the pit is broken down into benches first and then each of those is broken down into blasts or blocks. This leads to the schedule having hundreds or even thousands of database records containing information and data to be scheduled. Within each of those records, there is commonly a large range of information contained for each record, whether it is a spreadsheet or a specific scheduling tool, there will be rows of information. This information will cover areas

such as drilling, blasting, truck and shovel waste excavation, dragline excavation, ore excavation and ore processing.

I have certainly scheduled many mine plans, life of mine coal schedules in particular that might have something like the following:

1. 15 mine ramp areas;
2. 20 strips in each ramp area;
3. 7 blocks in each strip;
4. 4 horizons or benches in each block;
5. That leads to 8,400 records; and,
6. Within each record, there could easily be 200 rows of information.

In total, that leads to 1.7 million data points. That is a lot of possibilities for errors to creep in and a lot of data points to have to check for errors.

In mine schedules, those rows of data involve numerous calculations. For example, the conversion of the geological model ore volumes into processed ore tonnages, that is a lot of calculations that could go wrong.

Equipment Productivities

Every item of equipment that is scheduled requires production rate assumptions within the schedule. These assumptions could vary throughout the schedule for a number of reasons.

1. By block or location. For example, if the area is known to be a narrow face or shallow bench height, the digger productivity might be reduced.
2. By time period. For example, during the wet season, productivity might generally be lower.
3. By equipment method. If the equipment has different operating methods, then different productivity assumptions might be relevant. A great example of this are draglines which will have varying productivities when operating in different dig modes.

From one schedule to the next, parameters often change so, if time is available, checking that the right productivity assumptions have been used does make sense. In addition, there should be a routine process in place for checking and calibrating

productivities against historical results. By this I mean checking the assumed productivities against historical values and ensuring that your assumed productivity rates are realistic. For medium to long term schedules, this exercise should be carried out every time the schedule is run. For short term schedules, such as a weekly schedule, it is generally not necessary to carry out this exercise every time, but productivity assumptions should be checked at least once per quarter.

Time Usage Model

The Time Usage Model (“TUM”) is the planned hours worked by each item of equipment during each time period. The time usage model is critical because the production achieved in any time period is determined by the assumed productivity multiplied by the assumed operating hours. At most mine sites, TUMs have generally been developed to the point that they incorporate a significant number of delays. It is not uncommon for a TUM to have about 40 delay types in it, with those delay types generally divided into four categories.

1. Scheduled maintenance delay.
2. Unscheduled maintenance delay.
3. Scheduled process delay.
4. Unscheduled process delay.

Calendar

The calendar dictates the working days and hours, it is also used for some delays to dictate when they actually occur, for example, planned maintenance periods. A calendar can also be useful for creating different productivity rates during different periods, for example, reduced productivities during the wet season. Similar to the main database (underlying reserves discussed above), the calendar could potentially have a lot of records in it, with many rows of information for each record.

Face Position Updates

Before the schedule can be run, it needs to be calibrated so that the schedule is modelled from the correct starting points. That means every item of equipment has up to date face positions, so that they are starting from the right projected place when the schedule starts. This usually requires either a visual estimate of the equipment starting position or the survey team picking up actual face positions (in the case of short term schedules). However, for mid term and longer term schedules, it is typically more common that the schedule will need to be calibrated to start from the end of the short term schedule that the longer term schedule ties in with.

Equipment Sequences and Input of In Path Delays

Last, but by no means least, the sequence of blocks to be scheduled needs to be created or modified for every piece of equipment in the mine plan. For a large mine this can be a very onerous task and can be very time consuming. This task also requires a detailed knowledge of the mine, so as to understand the dependencies between activities and various mining areas. It is essential, particularly for shorter term schedules, that the mine planner knows the links between the activities in all the mining areas.

It is very easy to make mistakes in this part of the scheduling and it is just as easy for them to go unnoticed, particularly if there is not some form of reliable checking process in place. In addition to sequencing the activities, the mine planner will also need to add delays into the sequence where relevant. The most common of these delays is that of equipment deadheading from one mining area to another, which is not normally entered into the schedule via the TUM or the calendar, but is instead entered as a delay in the equipment sequencing path.

In addition to all of those commonplace items just discussed, we have piled more on our modern-day mine planners. Scheduling software now incorporates functions such as destination scheduling, haul route modelling and automated planning (requiring sequencing rules). We are continuing to add more and more for our mine planners to review and widen the scope for things to go wrong. I liken it to the widely known effect on equipment availability of continuing to add more technology to our mining fleets. Every time we add something new to our trucks

and diggers, such as weight measurement systems, fleet management systems or fatigue management platforms, we add something else that can break down and, therefore, lower the equipment availability.

With all of these data input areas requiring validation before commencing the schedule, it almost makes you wonder how the plan could not be full of errors.

Why Is “Bullshit In” A Problem?

A plan is only as good as the inputs to that plan. As I said at the start of this chapter, bullshit in, bullshit out. Without carrying out the correct due diligence on mine plan inputs, the probability of it being a bullshit plan increases dramatically. In my 30 years of scheduling, I have noticed that nearly every time we found a mistake in a schedule, it has been an error resulting in a downside to the plan rather than upside. By downside in the plan, I mean that it takes longer than the plan said it would and we likely won't achieve target. Now “nearly every time” might be an exaggeration, but certainly that's the way it has felt.

Let's be clear, by an error in the plan, I don't mean the production of an optimistic schedule as will be discussed in Chapter 13, where the error is in the judgment as to how long tasks might take. I mean an innocent mistake, such as where a value is incorrect in the database as it was calculated wrong, a number was inadvertently typed over a formula, or a block was accidentally left out of the schedule, etc.

Maybe this isn't just a coincidence? Let's investigate further and explore the types of innocent mistakes that occur, for example:

1. Forgetting to sequence a record, for example a block of excavation or maintenance delay.
2. Forgetting to enter a value in a record that is in the database. The record is there but contains no value for one or more of the parameters such as drill metres.
3. Forgetting an interaction between two activities, for example, forgetting that a block can't be drilled because the haul trucks are still going to be running through that block.
4. Accidentally deleting a record such as a block of excavation.

5. Accidentally zeroing out a value in a record such as drill meters.
6. Overwriting a formula in a record with a value.
7. Forgetting to update a face position.

An error such as any of the above is going to cause the schedule to be wrong and it can cause the schedule to be wrong in one of two directions. First, where the error will lead to the activity actually finishing ahead of plan, for example, having a scheduled quantity that is less than the quantity that is actually there. I describe this as an upside error, the error leads to an upside result in that the actual execution finishes ahead of plan. The alternate outcome to this is the downside error, where the activity takes longer than planned so the actual finishes behind plan. Downside errors lead to a negative result in that you don't achieve the schedule targets.

Now, the question is this, which is more likely, a downside error or an upside error? We can draw some conclusion on this by reviewing the list of seven potential innocent mistakes from above and mapping whether they might lead to upside or downside errors. This exercise has been carried out in Table 8 - 1 below.

Type of Error	Database (Underlying Data)	Time Usage Model	Calendar	Productivity Assumptions	Equipment Sequence
Forget to include a record	✗	✗	✗	✔	✗
Forget to enter a value in a record	✗	✗	✗	✔	✗
Forget an interaction between two activities					✗
Accidentally deleting a record	✗	✗	✗	✔	✗
Accidentally zeroing out a record	✗	✗	✗	✔	✗
Overwriting a formula in a record with a number	✔ Upside or Downside ✗				
Forget to update a face position					✔

- ✔ Creates Upside
- ✗ Creates Downside
- Not Applicable

Table 8 - 1 Schedule Error Mapping

As you can see from this table, of the 30 possible error results, 21 of them led to errors that have potential downside. That's a 70% probability of a downside error and that is assuming that each of the 30 possible scenarios has an equal probability of occurring, which is unlikely to be the case. Depending on what the mistake is, where it occurs in the schedule process and how large the error is, it can be impossible to reschedule the plan so as to still achieve targets, therefore costing the site money.

However, this is not the only issue with a lack of due diligence in mine plans, there is another significant pitfall and the consequences are potentially larger than that of not achieving targets. It goes to credibility, trust and confidence in the plan and planner, which is a critical element of successful mine planning and execution.

There is nothing like a mistake coming to light in very public circumstances to create a substantial credibility issue for the mine planning team. I've experienced it (see some of my stories in this book) and I have seen it happen to others many times. The credibility issue occurs for both the planner and the plan itself, even the most innocent of mistakes makes the planner look unprofessional. It is very difficult to come out the other side with your credibility intact and a lack of credibility breeds a lack of confidence in both the plan and the planner.

What Is The Solution?

Graduating as a civil engineer into the Department of Main Roads, I've always remembered the mantra, "there are three things that are important in road design – drainage, drainage and drainage". Maybe the same could be said of mine planning... "There are three steps to not looking like an idiot – check the plan, check the plan and check the plan". Invest time in checking the plan before you start, it is as simple as that. It is not possible to diligently check everything, but it is definitely worth a rigorous search for issues that may be material before scheduling.

Small mistakes and issues won't ever bring you undone or cause anything, but minor embarrassment. Large mistakes, however, are different. They are very embarrassing at the time and may stick with you well into the future. You will always be the person remembered for "that" particular mistake and you can become

infamous for all the wrong reasons – I’ve seen it happen. I would be putting some effort into ensuring that there are no major blunders in my schedule.

There are a number of methods for avoiding mistakes in schedules and these can be categorised as follows:

1. Tools.
2. Sanity checks.
3. Checklists.
4. Notes.
5. Peer review.
6. Accountability.

For all the above solutions, start first with high level analysis of the areas that will trip you up the most, that is where most errors occur that will cause the greatest amount of pain. In any schedule, there are critical areas, ensure you are aware of the critical areas in your schedule and design solutions to ensure that these areas are error free. For example, at many mine sites, drill and blast is always critical within the schedule, it usually has a short time frame in which it must occur and it needs to finish on time, otherwise excavation equipment parks up. So, drill and blast is an area to focus on and pay some attention to.

Tools

Let's start with the first item on the list above, tools that might help. Many scheduling software platforms allow you to create images or plots of some form. I have always found these useful as part of my due diligence process, it is typically my first step that I go to. If I turn up to a new mine site to carry out a schedule, which I have done many times, the first thing I do is see if they have block polygons built into the software, allowing me to generate a range of plots in plan.

For an open cut coal mine, here is a typical range of plots I create showing a plan view of the strips and blocks that will very quickly tell me if something is wrong in the underlying reserves:

1. Prime strip ratio, which is prime waste volume divided by total coal tonnage.

2. Total coal thickness.
3. Coal yield.
4. Number of dragline passes.
5. Total dragline rehandle.
6. Total profit per tonne.

As well as producing graphical plots in a plan view, some scheduling software allows other options, such as bar charts. These can be very useful, as they allow you to start at a macro level and then drill down. They can be helpful for finding spurious records by looking for spikes in the bar chart. For example, a check on drill metres could be carried out by creating a calculation of drill metres in a block divided by the block area, therefore providing an average drill hole area. Plotting those across all the records will show any spikes where there could be a drill metre calculation error.

If the schedule is in Microsoft Excel, or some other form of spreadsheet, then it is possible to use a mix of charting plus conditional formatting and filtering. Again, applying the similar logic of creating additional calculation columns or rows where required to help these checks to proceed.

Sanity Checks

One thing that I haven't seen enough of in mines that I've worked at is the simplicity of sanity checks. For me, sanity checks involve having an expectation of what it is that you will see from your schedule, it is that gut feel about what is right based on past experiences. Sanity checks are about removing the "black box" effect that occurs when you can't follow the software processes and therefore you have to just trust that the software is going to give you the right answer. For me, the only home in mine planning for black box analysis is in optimisation software. No mine plan should ever resemble a black box, the mine scheduler should always have a thorough understanding of their schedule.

Checklists

I would strongly recommend reading the book “The Checklist Manifesto” by Atul Gawande. This book primarily discusses the extensive use of checklists in the flight industry and how such checklists should be used in other industries. It details how the use of checklists have saved lives and have changed the medical industry. In my experience checklists are rare throughout the mining industry, and those I have seen in use were more about company governance processes rather than making a change for good.

Within the creation of a good mine plan would be a good place to start using checklists more. The Checklist Manifesto explains very well where to use checklists, where not to use them, the different types of checklists and the processes for developing, testing and implementing them. The book provides plenty of examples where checklists are critical, one example of this is when a routine process is regularly and frequently interrupted by people or other activities. This sounds like the weekly scheduler who is trying to put together a schedule and is frequently interrupted. I challenge you, what do you have to lose in creating a checklist? Go ahead, try one now.

Notes

Mine scheduling software have varying strengths; however, I haven't come across any scheduling software that allows good quality notes to be written in the schedule. I would say that this is an area that is fairly poor in general. Excel is better for this as you are free to write notes and comments throughout the spreadsheet. But this also opens us up to a secondary issue, whether the scheduler will actually take the time to write notes. So, we have two potential problems, either the scheduling software does not have the capability for notes to be left or, where it does, notes are not left as the scheduler doesn't have the desire or the time.

The writing of notes is underrated and should be applied more. There have been countless occasions I have looked at my own schedules a month or two later and couldn't remember why I did something, especially when it was outside of the norm. Notes are definitely required for these abnormal circumstances, but notes can also be made for the future, a reminder of something that should (or shouldn't) be done

within the schedule. For example, don't excavate Block A while drilling is occurring on Block B. It's almost like a checklist.

Keeping comprehensive notes will certainly help in avoiding mistakes, but they are even more critical if there is the potential need for someone else to work on the plan. As a consultant, I have regularly filled in for planning engineers in creating schedules and often I was called in late in the scheduling process. There are many times that I turned up to schedule a mine I had not previously worked at, so I had no background in how the mine worked. Notes would have been exceptionally useful.

One thing I noticed though, is that I was never engaged at a new mine site for the first time and given the task of creating a weekly schedule. Creating the schedule that is to be executed would be an extremely difficult task for someone unfamiliar with the site, and extremely risky for the mine site itself. If you're outsourcing some of your scheduling, my advice is to never get someone in to create an execution schedule.

Peer Review

Lastly, build in a peer review process of some form, an independent check of the plan, with an alternate person carrying out a high-level sanity check. I recommend that the reviewer uses a checklist that has been specifically designed for this purpose. The peer review could be by, as the name suggests, a peer of the planner. The process will work well if it's another planner who has a good understanding of the mine and the schedules. For example, a planner who is responsible for carrying out scheduling on a different time frame at the mine site.

However, my preference is to apply a model that I have called the "spruikers and the doers". The doer is the junior person at the mine site who actually carries out tasks such as drill designs, dragline designs, mine scheduling, etc. The spruiker is the more senior person at the site who is effectively the communicator. They are the facilitator who deals with the Production department who execute the plan, as communications are a critical component in the mine planning process. The mine plan is a story, it's a story that you want people to take on, understand and implement. That story should be told by an experienced person, experienced not only at the mine site, but experienced in interactions with production personnel.

The Production team understanding the mine plan is of critical importance and this process works by the doers carrying out the work to create the plan, before the spruiker uses their experience and effective communication skills to help the Production team understand the plan. The spruiker is the person who carries out the peer review process, they check over the plan and then they take on ownership of the plan before they sell it. The spruiker then communicates the plan, monitors the plan and ensures that the plan is executed. The doer continues to be involved, they modify the plan when required and they learn by watching the spruiker, so it becomes a personal development process as well.

9

CRIME 6: Beware Of The Enablers!

Many years ago, when I was working in an open cut coal mine, I had a great opportunity to observe interactions between the Mine Planning department and the Mine Production department. And the cost of failure.

It was a standard strip coal mine with a dragline and we had a low spot at the end of a strip that had filled up with water following rain events - there was approximately 20 metres depth of water in it. The Technical Services Department (mine planning) started flagging early the need to dewater this area before we blasted the dragline waste that was sitting above the water. We started flagging this while we still had access to the area, so that the mine services crew could install pumps.

It didn't seem much of a priority with the Production Department and our requests fell on deaf ears. Our plans were ignored and so it came to the point where the dragline crossed the coal haulage ramp, which meant that access to the area for installing and maintaining pumps was blocked. With access lost via the coal haulage ramp, this meant pump access through the dump spoil had to be pushed in by dozers. The issue became more pressing as it became obvious that there was going to be a problem.

Week after week in mine planning meetings, our Technical Services team highlighted that this was a critical issue. After a period with no success, we began quantifying the rate at which the water needed to be pumped out and therefore the number of pumps required. Obviously, as time went by and pumping never began, that rate just grew each week.

The water wasn't pumped out, although there were some token efforts, but as the dragline approached this mining area, we still had a pit with about 20m depth of water in it. We got to the point where we had no choice, the dragline was about to park up waiting for the blast to be fired, so we fired the overburden blast into about 20m of water. This did make for spectacular viewing as it sent a tidal wave of water up and over the low wall spoil. But as a result of blasting overburden into water, the material had no strength, it was like mud. Once the dragline began excavation in this area, the consequences of the lack of dewatering became very clear.

A necessary part of dragline operations is for the dragline to build a bridge (made from waste) out into the void so that it can excavate the blasted spoil away from the coal edge and place it into position in the final spoil. The dragline couldn't build a

bridge that would stand up due to the low material strength, it just continued to crack and fail. We reached a point where the dragline couldn't carry out its normal excavation method, which of course was a huge emergency. This required the Technical Services team to come up with a plan to allow dragline operations to continue.

The new design for that plan was that we had no choice but to walk the dragline around into the lowwall spoil. The dragline had to excavate a path through the spoil, so that it could walk out to the lowwall edge and rehandle the mud out of the pit and dump it into the lowwall spoil. To dump it into the lowwall spoil, the dragline had to first excavate a dam holding area to hold the mud. The dragline then had to rehandle dry spoil from the spoil area and place it back into the cleaned-out pit so it could be used to build the bridge required. Needless to say, this whole exercise cost the dragline at least a month in production time. This was a month not spent uncovering coal, which led to demurrage for ships sitting off the coast waiting to be loaded. It was an extremely costly exercise.

When I reflect on this event, there are two primary issues that strike me. The first is the primary reason for it occurring in the first place. That is that the team who were charged with executing the plan were not the owners of the resources that were required for execution of the plan. The mine services crew were the team charged with dewatering this area, however, this crew did not have any dozers that were allocated to them. Consequently, the Mine Services Superintendent was continually trying to get dozers from other Production Superintendents. The other superintendents have their own priorities, therefore, the dozers were never made available for pushing the low wall track down to access the water. The significant misunderstanding here was that drainage and water management are so often an enabler for excavation, which I have also seen on numerous other occasions.

The other interesting thing that I noticed after the whole event happened, is that somehow it became a Technical Services stuff up. We wore the blame for the whole debacle. My guess is that because we came up with a solution that was so different to the norm (out of necessity), it was easy to shift the blame and suggest it must have been a poor design by our team in the first place.

What Is The Problem?

Do you have the situation where blasts at your mine site are consistently fired late? They are fired late because the drilling ran late, the drilling ran late because the drill bench preparation was not completed. Or, even worse, it was completed but was inadequate because it was rushed and subsequently it requires preparation again. This often happens because the dozer or the grader required for the bench preparation was assigned to other (theoretically) higher priorities.

Or do you have the situation where the excavator is planned to start digging in an area, but can't because the road wasn't pushed in for the pump crew to dewater this area allowing the excavation to start? Again, this commonly occurs because the dozer required by the pump crew had other seemingly higher priorities at the time. If so, then your mine, like countless other mine sites I have come across, is suffering from a lack of attention to enabling activities.

Every mine site has enabling activities, these are critical items that must be completed before another activity can start. They "enable" that activity and they have four key features:

1. They happen immediately before a key production activity, such as drilling or excavation.
2. They're not scheduled. But it is bigger than that, not only are they not scheduled, but they generally don't have base task quantities calculated, which allow estimated times to be calculated. Therefore, they have no timeline.
3. They are not scheduled so they are not monitored against a plan. Therefore, you don't know if you are ahead or behind plan on this key enabling activity.
4. They're not a key production task (or apparently so). So they are not prioritized and this is not helped by the fact that they had not been scheduled.

It is completely natural for enabling activities to be a case of "out of sight, out of mind". Drill bench preparation is one of the best examples that I've come across of a key enabling activity. I have worked at numerous mine sites where it has been a constant complaint and it perfectly fits all of the four criteria listed above.

In the case of drill bench preparation, it usually has another fifth feature which exacerbates the issue. The person held accountable for an enabling activity does not have authority or responsibility for the equipment or personnel required to complete this task.

This is generally because the particular enabling activity is not seen as a key part of this person's role and is a sporadic task, not carried out on a full-time basis. It may not even be included in their job description. For a Drill and Blast Superintendent, they are in charge of the drills and the blast crew, as that is their core role. They are required to drill metres, load explosives, and fire shots. So, generally they do not have any dozers, graders, water trucks, or other ancillary equipment under their control. Those equipment items are under the control of other Superintendent roles at the mine site.

Those Superintendents have their own tasks, KPI's and objectives that they are held accountable to. When the Drill and Blast Superintendent requires any of this ancillary equipment for bench preparation, they are brought undone by the four features of enabling activities outlined above. It is not perceived as a priority task, certainly not if the transfer of that equipment on to bench preparation is going to impact the production rate of an excavation activity. The argument for getting access to the ancillary equipment is not helped if the task hasn't been scheduled so it is not possible to convey a time frame of when the equipment is required by, or the urgency of the requirement. If the enabling activity hasn't been scheduled, then the priority can't even be adjudicated by an independent person such as the mine planner.

There are other enabling activities, these all have the same features, some examples are as follows:

- Attain access to an area, particularly for a mine services team (pump crew). When access is required for a truck shovel excavation, it happens very quickly. When access is required for anything else such as the mine surfaces crew, it is a very low priority.
- Excavation of drainage sumps, often drainage is a prerequisite before any mining activity can occur in an area that is subject to drainage issues.

- Top of coal clean up and inspection. This is typically not as big an issue as the other examples, because it is seen as a priority, but it has the same features as those other enabling activities.

This is just a small set, every mine is different and the enabling activities at your mine site might include other tasks. The easiest way to find the enabling activities at your mine site is to look for the issues that most commonly cause delays. They are often the subject of much discussion, especially at planning meetings and there is generally evidence of frustration, particularly by those responsible for the next task downstream in the process.

Why Is Lack of Consideration Of Enablers A Problem?

The best mine plan in the world becomes completely irrelevant if other issues lead to the plan being wrong. If enabling activities are ignored within the mine planning process, they will invariably hold up critical activities. Holding up critical activities will result in scheduling targets not being achieved. Or, if the activity being held up is not critical, then there is the potential to turn that non-critical activity into a critical activity, meaning that you've just increased the risk across your entire mine plan. The lack of consideration of enabling activities is a critical issue, you are effectively missing critical interactions within the schedule and therefore the mine schedule won't work.

The other issue which occurs is that it leads to volatility in resource requirements. A task becomes urgent so resources need to be thrown at it, but those resources may actually have been required for other planned tasks, which consequently turn into urgent tasks themselves. This results in schedule disruption beyond this one activity and, in fact, it can impact other scheduled tasks as well. It leads to an inefficient use of resources and it leads to tension and frustration between Supervisors and Superintendents, as well as a culture at the mine site that is less than it could be. Talk about a vicious circle.

What Is The Solution To Enablers?

I'm sure you've heard that often quoted definition of insanity being “continuing to do the same thing and expecting a different result” (usually falsely attributed to Albert Einstein). Ignoring enabling activities within the scheduling process definitely meets that definition of insanity. To solve this issue, start with determining what the enabling activities are within your mine plans. If it's not immediately clear what they are, create a process of logging why each schedule was held up. This will highlight the causes for schedule delays and then check whether the four features of enabling activities discussed earlier in this chapter apply.

If there is a frequently occurring enabling activity, such as drill bench preparation, then the answer is simple, schedule that activity within your mine plan, it is as simple as that. This discussion relates to short term schedules, it's not required in longer term schedules. In longer term schedules, it is sufficient to just leave a time gap between tasks either side of the enabling activity. But, in short term schedules it's definitely required, they are important enough to schedule and then monitor that activity.

The other advantage that stems from scheduling an activity is the monitoring. A scheduled task can be monitored so it is possible to measure whether it is ahead or behind plan. If the activity is not routine, for example digging a sump in a particular pit area, then there is no need to schedule it. However, there is still a need to carry out a calculation of the estimated quantities to allow the estimated activity time to be determined. Following this step, it is possible to recognise the date it is required by and discuss this at planning meetings and highlight as required, so it does become a critical activity and therefore equipment can be prioritised. If this process isn't carried out, then you're in danger of meeting the definition of insanity.

It is relevant at this point to compare mine scheduling to project scheduling. By project scheduling I mean those projects usually scheduled in Microsoft Project, Primavera or another project management software platform. Construction projects differ from a mining project in that they will have primary production activities like formwork setup or concrete pours for example, but they also then schedule every other task that requires completion, not just the main production activities.

Why is it that mine sites have a comprehensive scheduling process that is carried out on a very detailed and repetitive basis in complex software, but then we don't schedule everything? We just selectively choose some primary production activities.

Another crucial difference is it that a key functionality built into project management software such as Microsoft Project is the use of predecessors and successors for tasks. These were explained in Chapter 2 (Terminology) in the discussion on dependencies and interdependencies. Tasks are linked together, if the timing of the predecessor changes, then the timing of the successor task will vary as well. But in the mining industry, often our mine scheduling software doesn't incorporate this feature. We don't nominate the predecessors and successors, we actually just rely on the scheduler to remember this.

The following story provides another reason why and how you should use enabling activities. This story follows on directly from the broken stocks saga in Chapter 1.

For me, the learning from the broken stock saga was one of the very thought provoking concepts to come from this period in my mining career. In theory, the Short Term Planning Department were responsible for planning the time period from zero to three months; however, in reality, the Grade Controllers were actually the VERY short term planners, even though they were positioned in the Mine Production Department. The Grade Controllers dictated where excavators were located each shift and which of them were operated, as they were tasked with producing specification quality stockpiles. I would argue that this is really what **everyone** is at the mine site to help facilitate – producing the required quantity of saleable product.

The Grade Controller dictated where all mining equipment worked and the Shift Supervisor was charged with maximising equipment productivity, given the equipment working locations. Consequently, why should the Mine Planning Department even try and schedule where the shovels would dig on a short term basis? The reality is that they don't dictate this, the Grade Controller always will, as they carry responsibility for hitting the mine quality and quantity targets.

But there is a way the mine planners can lead the Grade Controllers where they want them to go, which is to follow the three month plan. They can indirectly facilitate the shovels working where they want them to, without scheduling them, by

focussing on scheduling of the activity that is the enabler for excavation. In most instances this is the mining activity immediately prior to excavation in the mining sequence, which is usually drill and blast. If you focus on scheduling the drilling and blasting, which absolutely should incorporate the management of broken stocks (as per the broken stocks saga), the shovels are sure to follow. They have to as they have no choice, if it is not broken they can't dig it.

The Short Term Planning Department continually provided feedback to us that we (Grade Control) were digging the wrong pits and weren't conforming with Annual Plan pit quantities. Now I understand the need to follow longer term plans, excavating differing quantities in pits from that planned over a sustained period can have serious implications on numerous elements, such as maintaining sufficient mining faces, infrastructure and capital requirements. But as the Grade Control Superintendent, if I'm forced into a choice between producing the correct quality product, or following the Annual Plan, it's a no brainer – the quality product wins every time. And when you have eight quality targets and broken stocks are off specification, then you have very little flexibility, as there are few combinations of available ore blocks that achieve target.

But if Short Term Planning Department focussed on where to drill, then they could follow the Annual Plan pit quantities, as long as broken stocks were maintained within target bands. If it's not possible to do this, then that indicates one of two primary scenarios. First, Annual Planning is normally carried out at a granular level and typically on a monthly time frame, which can disguise the greater variability that occurs when scheduled on a weekly basis. Or second, that it was a bad Annual Plan in the first place.

The irony in this, is that after 30 years of mine planning, I would say that the drills are probably one of the easiest items of equipment for Mine Planning/Technical Services Departments to schedule and actually have some control over. Drill and Blast Superintendents and Supervisors generally don't care too much where they are drilling, as long as they are productive because you're not constantly moving drills around or firing very small shots. So give them a workable plan for their drills and there is a high probability they will follow it. Truck/Shovel Superintendents and Supervisors on the other hand are often more particular about where they will dig so, I would suggest at many mine sites, it would be much more difficult for the Planning Department to have some control over those trucks/shovels.

Enabling activities in actual fact present an opportunity, they may not be the true bottleneck for the mine site, but they are their own mini-bottleneck and lead to the potential for a Theory of Constraints (TOC) approach to be used (as further discussed in Chapter 12). There is the potential for the enabling activity to be positioned as a constraint and then to utilise that constraint as per TOC. I would suggest after reading Chapter 12, come back and review this chapter as to how you would leverage off the application of TOC and enablers together.

What follows is an example of an enabling activity that I believe is under-utilised by mining companies. For the mining industry, inventories play a critical role in allowing us to manage variability in our plans. In that scenario, there is at least one additional inventory we should be managing, that is “**Available To Drill**” inventory. I haven’t come across it in use at any mine sites, but I believe it could prove to be essential to efficient operations.

Many mine sites measure and report against blasted inventory, drilled inventory or both. But I haven’t come across any sites completing the cycle by scheduling, measuring and reporting “available to be drilled” inventory. At many mine sites, the drills do have lower utilisation rates, subsequently excess drilling capacity (on a short term basis) is often required to achieve annual drill metre requirements. Because on many occasions, drills are parked up as a result of there being nowhere for them to drill.

It is the third piece of the pie that is missing. There are three steps to having sufficient quantity and quality broken stocks – having the right places to put the drills, drilling at an adequate rate and then blasting at an adequate rate. Across the three it is possible to determine a total inventory range which will ensure that a mine site maintains adequate broken stocks (blasted inventory). The three inventories can be combined on a stacked bar basis per digging unit, as shown in Figure 9 - 1, to provide an immediate picture of how things are going, where the current issues are and where the future issues may be coming.

Immediate issues occur where there are insufficient broken stocks, so Shovel 12 above has only about 3 days digging, that’s an immediate problem. However, it does have plenty of drilled and available to drill inventory though, so the longer term looks satisfactory. Shovel 11 on the other hand is a significant issue in the long

term, it has 13 days of available dig material, but no drilled inventory and worse still, nowhere for the drill to even drill.

What is extremely beneficial in this process is that each inventory is an early indicator to the following inventory. For example, if the available to drill inventory is low, then unless the drilled inventory is very high, there will be an issue with drilled inventory in the future.

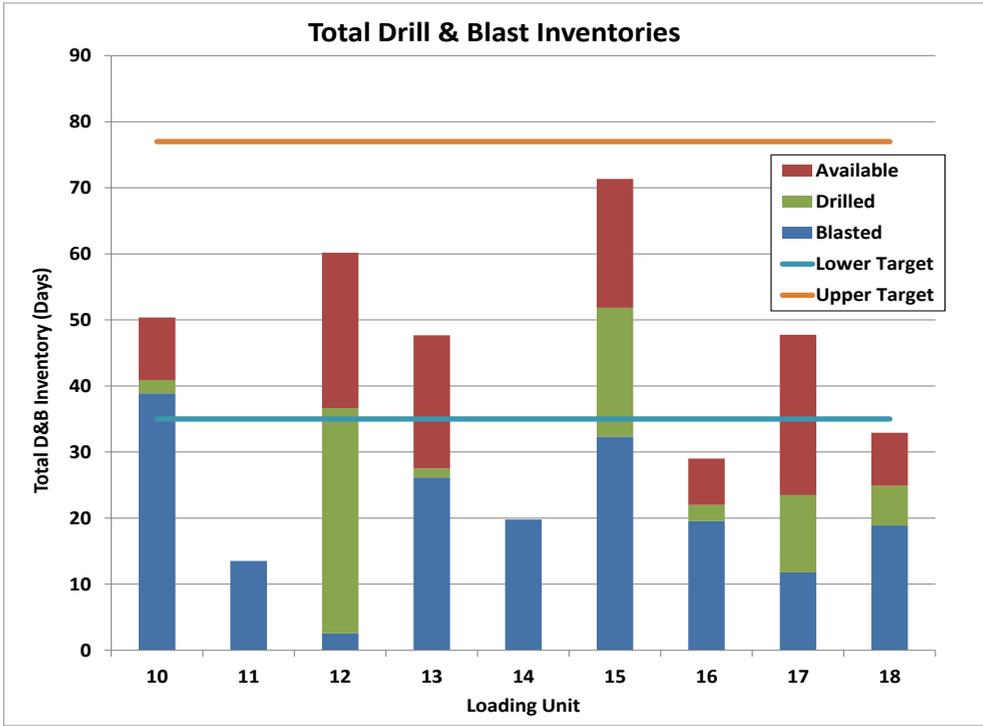


Figure 9 - 1 Total Drill & Blast Inventories

At mine sites where excavators move between mining areas regularly, this process could be carried out on a pit (or bench) basis instead. It will serve the same purpose. The aim is that the regular use of the three inventories will highlight problem areas in time to see issues looming and regulate them. Who can't argue that the earlier you see problems coming, the better that is for any mine site.

10

CRIME 7: A Law Unto Themselves

I once worked with a Production Superintendent who liked a “robust discussion” (which is another way of describing a good argument). Not only did he like an argument, but he prided himself on winning them. I was finding this a real struggle at the time, as I’m a great believer of the Senior Planning Engineer and Production Superintendents having a good working relationship and combining well to get everyone else on side.

Fortunately, my manager at the time was one of the best mentors I had and he gave me some sage advice. That was not to worry about winning the small battles, but to concentrate on winning the large battles, the important decisions that needed to be won. And that my adversary didn’t judge arguments won on the importance of the argument, but purely judged performance on the number of arguments won. So, if he won eight out of our ten battles, he was happy. If I won the two battles that needed to be won, I was happy, and this process worked remarkably well. Subsequently, I just fought hard and stood my ground on the battles that needed winning and let him win the bulk of the smaller battles. I’d even trade-off one disagreement I wasn’t so concerned about for one that I was concerned about.

What Is The Problem?

It’s not uncommon for the Production Team to have either no faith at all, or very little faith in either (or both) of the mine plan or the Mine Planning Team. This is commonly due to a history of poor plans being produced or, as far as the Production Team were concerned, they were poor plans. One of my favourite sayings is “perception is reality”. If the Production team perceive the plans as poor, then to them that is real. In the majority of cases, they perceive the mine plan as a poor plan because it doesn’t last for long before it is wrong (a short shelf life). And, ironically, as discussed in Chapters 4 and 5, that is often caused by high variability and the resulting slippage, which is something not under the Mine Planning team’s control.

Short shelf life mine plans lead to either equipment idle time or equipment relocation. From a Production team perspective, you are messing with their KPI’s. As soon as there is downtime of any type, they’re not happy, particularly if their bonus is tied into their KPI’s. Worse still is when the plan becomes too far from reality and so there is a need to re-run the plan. This constant generation of new

plans just reinforces what Production already think, which is that the mine planning is ineffective. When this happens on a regular basis, the Production team invariably arrive at a point of thinking to themselves “Why bother following the plan? I could do just as well by guessing it myself.”

As discussed in this book across a range of chapters, there are a lot of factors that lead to mine plans being wrong, however, those with the greatest impact are optimistic inputs and assumptions, and the lack of variability built into planning. But another issue that I've seen a number of times is a Production team or the Production Manager with too big of an ego. Thinking they know the operation better than the Mine Planning team and they can run it fine, they don't need some engineer telling them what to do.

This discussion relates to the execution plan, as the execution plan is the only one that the Production team actually have to worry about implementing. Although, I have come across mine sites that have measures in place regarding how well the Annual Plan is followed and hassle the Production team about not following the Annual Plan, or some other long term plan. In my opinion, hassle is the appropriate word to use, as trying to get a Production team to follow a mid or long term plan such as an Annual Plan is ridiculous. It is a completely pointless process, as they're always going to make short term decisions that suit today's operations, which is about keeping equipment working or shipping ore, rather than following the Annual Plan.

This story from when I was the Grade Control Superintendent at Paraburdoo iron ore mine and we had to build a stockpile every week with eight quality targets is a great example. The Mine Planning team were trying to hold us accountable to follow the Annual Plan. So they were measuring the percentage of material that we excavated from each mining area and comparing it to the Annual Plan percentage of material that was supposed to be excavated and trying to enforce that we should follow the Annual Plan better. But if you only have a limited choice of broken stocks, and the blasted material in a pit that is scheduled within the Annual Plan is not going to get your product stockpile to the right quality, you are not going to sacrifice the short term quality of that blended stockpile just to meet an Annual Plan.

I have a strong background in mine planning and believe in the plan. But to be honest, I did not take into consideration at any point in time the requirement to

follow the Annual Plan, when I was trying to build a production stockpile which already had so many constraints. Ironically, prior to my role as Grade Control Superintendent, I was the Senior Mining Engineer in the mid term planning team and was trying to enforce the need to follow the Annual Plan on the previous Grade Control Superintendent. Sometimes you have to look at issues from both sides of the fence.

Why Is Not Following The Plan An Issue?

Why is it an issue when the Production team decide that they don't need to follow the mine plan? Hopefully, you've read Chapter 3 and have a picture of what the world would look like in the mining industry with no mine planning, that should answer the question.

The Production team won't understand all the inclusions within the planning process. For example, the sources of where the information comes from, and the breadth of that information, which includes production, maintenance, processing, geology, geotechnical, etc. Not only that, but there are a wide range of outputs from the mine plan that numerous personnel in varying roles at the mine site rely upon.

The Production team are likely to focus on one or two items in the plan that they disagree with which may not be done so well and then use those as a reason to ignore the whole plan. This is an issue because there is significant complexity at a mine site, given the massive number of tasks, dependencies, interactions and all the data that's built in to a plan. Regardless of what the Production team think, they won't get it right nor do a better job than a well-planned mine schedule. Sure, there's a possibility that parts of the plan might be better than what is currently being produced by the Planning department. But overall, it won't be an optimal plan and, sometimes, they will get it really wrong and it will be a disaster.

It also breeds a very dangerous culture within the mine site. First, it negates the value of the technical work involved in mine planning. Secondly, this culture will impact other areas, if a team stops following the mine plan, then ultimately they may choose to stop following other guidance as well, such as mine designs. Or from there, choose to stop following other plans, such as the maintenance plan or processing plan. Other teams on site might mimic the Production team, if they don't

need to follow the plan then we don't either. The current mine plans may not be up to scratch but, in summary, it works better for the mine site as a whole if Mine Production decide to just follow the plan.

The Solution To Following The Plan

Production have lost confidence in the Mine Planning team, therefore, the onus is on the Mine Planning team to get that confidence back. To do that, the Planning team have to get introspective and start with why: why has the Production team lost confidence in the plan? I suggest creating a list of reasons why. The most effective way to do this is to put yourselves in the Production team's shoes. I'm a big fan of the mantra "walk a mile in someone else's shoes". So do that, pretend you're the Production team and make a list. In this process you have to get vulnerable and start with the mindset that the mine plan is inadequate. It would be a very challenging exercise (and enlightening) to find all the problems with the plan.

My belief is that this is where you must start because your first job is to win back the confidence of the Production team. You have to get them back on board with you and you won't do this if you don't start accepting responsibility. If you're making excuses, then you're not winning them over, so I would suggest, don't even start there yet. As soon as there is some analysis showing where Mine Production are contributing to the problem, you won't be able to help yourself and you'll point out what they need to fix as well. Then you'll start horse-trading your improvement for their improvement, which won't contribute to getting them on the same journey as you and giving them any confidence in your mine plans, or that you're doing anything about improving them.

If there is one thing I've tried to focus on later in life (as I've hopefully become wiser), it's to concern myself only with what I can control and, more importantly, not to concern myself with what I can't control. As the Mine Planning Manager or Superintendent, I can control mine planning, but I can't control what the Production team do with those plans. At best, I can only influence what the Production team do with them. So my suggestion would be to apply laser focus to what you can control.

Once you've shown your vulnerability and desire to improve your own area and the fact that it's unconditional (i.e. you're not attacking the Production team about their performance) then that should influence change in the Production team. They will feel obliged to help and will be much more amenable to contributing in a constructive way. Having identified a list of issues, you then need to work out a draft action plan and solutions and then discuss these with the Production team and get their input. That's about getting them on the same journey as you and assessing what to tackle first.

My preferred method for ranking the issues in the order in which to tackle them is through analysing them on the basis of effort versus return, Figure 10 - 1 shows an effort - return matrix where I have created labels to describe each of the quadrants. This methodology is also known as an Eisenhower Matrix, with the axes normally represented by urgency and importance, instead of effort and return.

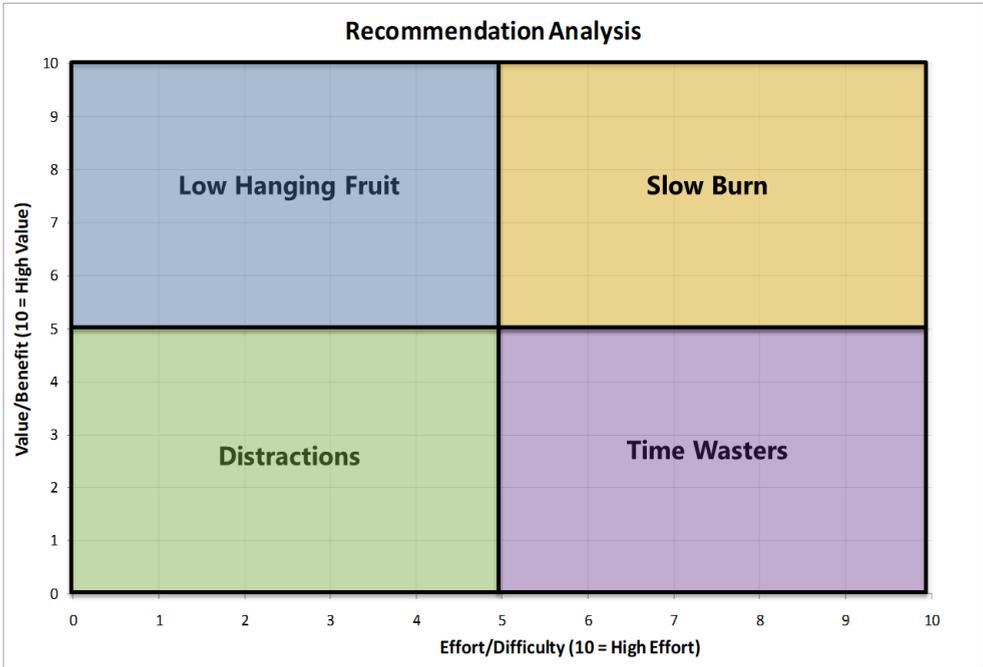


Figure 10 - 1 Decision Matrix

The explanation behind each of the quadrants is as follows.

- Distractions – because these issues only require low effort, it is tempting to tackle them early in the process. But don't, they provide little return and so are just a distraction early in the process.
- Time Wasters – these issues are high effort for little return, don't even go there, leave them alone as there will always be a better investment of time.
- Low Hanging Fruit – these issues are low effort and high return. This is where you start to get some quick runs on the board.
- Slow Burn – these are the high effort issues that can yield very high returns, an example of this would be switching away from current mine planning software to a stochastic scheduling tool (if it existed). This is the second quadrant to tackle as there will be some high yield projects here.

I've included in Figure 10 - 2 an example effort - return matrix resulting from an audit I carried out of a Technical Services Department. Note that as the allocation of scores for both effort and return are subjective decisions, I recommend this process is carried out by more than one person or is vetted by peers. I suggest starting with some "low hanging fruit" items first, those are the issues with low effort and high reward and are all about getting some runs on the board – some early successes. Getting some runs on the board improves people's confidence in the process and helps to drive change and should help with convincing Production to assist in the process.

I hope that from reading this book, including the other chapters, that you have built a list of ideas that you need to implement that will drive change and will help to keep the Mine Production team's confidence in the plan high.

But another approach is to review the primary causes of mine plans with only a short shelf life and then tackle those issues. When a plan is not very old before it has to be replaced or becomes irrelevant, that leads to people thinking "what is the point in having a mine plan?" The critical issue here is that most people (Production team included) don't understand the degree of variability in a plan and the impact that the variability has on both the accuracy and precision of the plan and its shelf life. They don't understand that the variability they themselves in the Production team have a hand

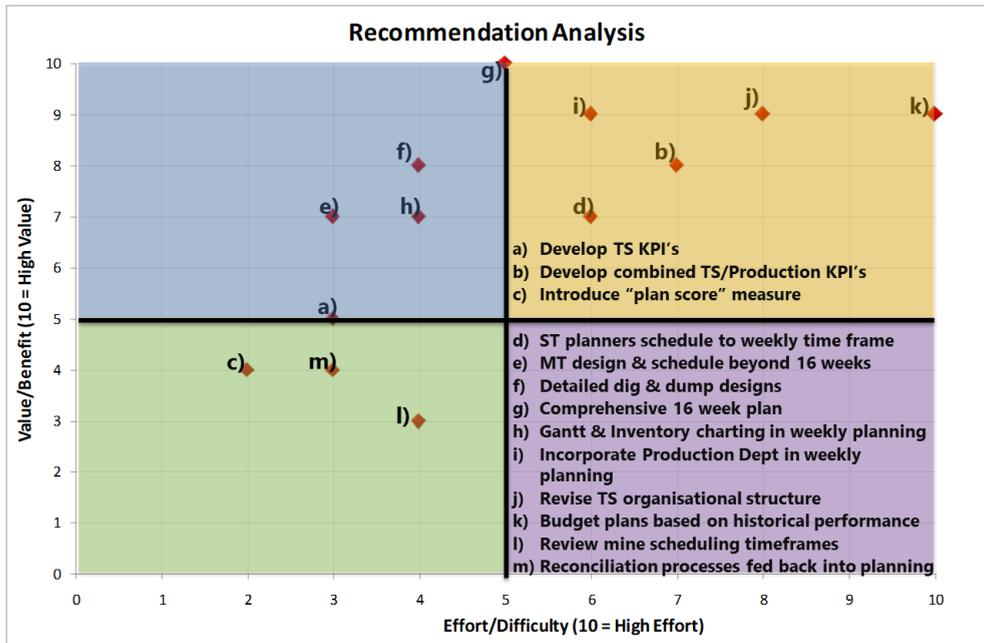


Figure 10 - 2 Technical Services (TS) Example

in creating, is one of the biggest causes for short shelf life mine plans. So, help the Production team understand that variability is a key issue, then they can actually help in the process of extending the shelf life of mine plans by contributing to the process of reducing variability.

The other primary contributor to short shelf life or inaccurate plans is optimistic input assumptions in the plan. Again, Mine Production have a key role to play in this issue, as they want stretch targets built into the planning so they can use it to motivate and push their teams. But using inputs that are known to be incorrect, as I have often seen, is a recipe for disaster. It guarantees that the mine plan will only have a short shelf life, as will be discussed in greater detail in Chapter 13. Help the Mine Production team to understand that the plans are never going to be right if they are used to manage performance. Therefore, the Production team, in wanting the plans to have optimistic inputs and stretch targets built into them, are actually negatively impacting the shelf life of the mine plans.

The strategy from there is simple, work out what is making the plans wrong, point that out to people and then ask them if they will work with you to fix it. If the answer is no, then the plan will always be wrong, but then that is not the fault of the Mine Planning team as the mine site has made that choice.

What follows is a related story, one that is not related so much to variability or optimistic assumptions, but an example of how organisational structure can be used to aid in the Production team following the plan. When I worked as the Grade Control Superintendent at Paraburdoo iron ore mine in the mid 1990's, this role was actually located within the Mine Production Team. I had a team of 4 Grade Controllers on continuous shift roster who were in charge of scheduling the loading units for the next week, with the aim of building a lump and fines stockpile to the target quality specifications. The Grade Controller on each shift was a peer of the Production Supervisor, they were equals and each reported through different Superintendents, but both reported through to the Mine Production Manager.

I remember at the time that my manager, the Mine Production Manager, held the Production Superintendent and his team, including the Production Supervisors, accountable for the quality achieved, i.e. whether we hit targets for the stockpile construction. I thought that was crazy, how can you hold someone in another team accountable when it was my team who were carrying out the planning of which ore we should mine and when? But I understand now, this was necessary to get the Production Supervisor to take notice of the Grade Controllers directions regarding where we should dig.

Disagreements on shift between the Grade Controller and the Production Supervisor were very common, primarily when the Grade Controller wanted to shut down a high productivity shovel and instead walk a front end loader into an unproductive dig face. That was because the Production Supervisor's primary focus was productivity and tonnes. But the Production Supervisor pretty much always followed the Grade Controller's directions, they had to, the ore quality was very complex and they had insufficient knowledge of the schedule required to achieve the quality targets. If the Production Supervisor was not held accountable for achieving the plan qualities, then when it involved an unproductive dig, they would have ignored the Grade Controller's request.

Here's another story about how to work with the Production Team to keep them on side. When I was Technical Services Manager at Burton Coal Mine, I was working for Thiess Contractors, who were operating the whole mine. Thiess have a very strong culture around productivity and profitability and one way they cultivate this is through the personnel they employ in Production Superintendent and Production Manager roles.

Unlike large mining organisations, who will often move mining engineers through these roles on their way to executive level management, Thiess typically follow a different strategy. Production Superintendents and Managers were nearly always ex operators who had moved through Production Supervisor roles. This was key to maximising equipment productivity as the Superintendents and Managers could coach and mentor from experience and identify issues that personnel without operational experience couldn't.

So I worked together with a Mine Production Manager who fitted the bill exactly for what Thiess aimed for in their operations leadership. He was an ex operator, about ten to fifteen years my senior, extremely focussed on production and making money for Thiess and not keen on taking advice or guidance from me. It didn't take me too long to understand that I was going to have to work out how we were going to work together, or things would not go well.

We held a number of discussions regarding what our roles were, what was important to each of us and how we could each best support each other. From those discussions, we agreed on the following.

- Technical Services was a service team and Production were our primary customer, we were there to provide the services as required by Production to ensure that the mine achieved all quantity, quality and financial targets.
- It was not Technical Services role to tell Production how to set up and operate the pit. So all execution decisions, such as which equipment to operate, which dumps to haul to, how to set up dig faces and which way to drill blast patterns were all decided by Production.
- It was not the Production role to tell Technical Services how to design and plan. Only Technical Services had the overall understanding and big picture of the mine over the weekly and beyond time frame so as to understand the

best decisions in this space. But that this work would be carried out on a consultative basis between both teams.

- That Technical Services were accountable for the quality of the plan, Production were accountable for the quality of the execution of the plan and meeting plan targets.

It worked brilliantly, we had a very clear set of boundaries around what our roles were and as long as we all worked within those boundaries, the mine worked efficiently and smoothly. Better still, the Production Manager and I maintained a great relationship for the three and half years we worked together.

11

CRIME 8: Are You Scheduling The Right Activities?

A couple of years into my mining career, I had a surveyor working for me and one of his favourite sayings was “a little bit of inaccuracy, saves a lot of time”. It was funny, because at the time I thought that was somewhat contradictory to the role of a surveyor and I actually didn't agree with his philosophy. It took me a fair few years, and in particular a number of years of working with and supervising mining engineers who were very carried away with details, for me to come around to that surveyor's way of thinking.

My mining engineers were spending a lot of time when carrying out long term designs or schedules in getting things exactly right, to the decimal point. Subsequently I was spending time continually trying to reiterate to them that they needed to consider the level of accuracy required for the tasks they were working on. And that maybe a long term design didn't need to have so much time put into it being exact, because the issue with long term designs, is that they're never actually going to be built. It's the short term design that is going to be executed, and therefore needs to be correct.

This has caused me to reflect on accuracy requirements and how in actual fact, I think that surveyor was right. If you think about surveying in the mining industry, it works in large scale. It's not a matter of building a high rise building or a highway, where there is a need to be accurate to the millimetre. So, in actual fact, a little bit of inaccuracy was worth the time saving for him.

What Is The Problem?

It's always important to be aware of the level of detail required for the mining engineering activity that is being carried out and the primary driver of the detail required is the time frame. Is it a short term schedule/design, or a long term schedule/design?

Does any of the following sound familiar? You don't schedule drill and blast, which is a critical enabler for excavation tasks and often a sprint activity. Maybe within your schedule you just leave a gap between activities to allow for the time that drill and blast takes? Or maybe there's no recognition of it at all in the schedule? Maybe you don't schedule other enablers, such as drill bench preparation and instead just insert a gap between tasks, similar to what was done for drill and blast? Or there's

no recognition of it at all? Or conversely, maybe you schedule too much equipment. For example, at a life of mine level, you individually schedule all of the digging equipment.

As an aside, although I haven't really discussed them in this book, at a life of mine level, I'm a big fan of auto schedulers and optimizers. I don't believe in them at short term planning levels, but at long term levels they really do serve a function. That comes back to the purpose of the plan, anything beyond the shortest term plan is never going to be executed. A life of mine plan is primarily for information and decision making and to achieve this purpose adequately, you don't need to spend hours on detailed planning. This can be done with the use of an auto scheduler and setting up rules for it to schedule by. Yes, the rules can take a long time to set up initially, but then once they're created, running the auto scheduler is a relatively quick process and can allow a number of iterations to be run. To run schedules in future years, it is just a case of reviewing the reserves, and the scheduling rules, and then hitting the go button on the scheduler. This has the added bonus that the engineering time previously spent on scheduling, can instead be put into strategy, trade-off studies or other value adding work.

Over the last twenty years, scheduling software providers have introduced two new functionalities that have become widely used and almost universal to be carried out when scheduling in some companies. These are destination scheduling and truck haulage models. They are useful functions, but save them for when they are absolutely required and I would strongly recommend against running them every schedule, as I have seen at some mines.

Destination scheduling involves scheduling the dumping of the waste in addition to scheduling the excavation, so you have doubled your workload. As you can imagine, there is an inordinate amount of time required to create the dump reserves, set them up in the schedule, and then sequence them. So you have dramatically increased your time to schedule and to be honest, I believe they are a gigantic waste of time.

Destination scheduling should typically be carried out for two reasons. Firstly, to check that the spoil fits and that there is a positive spoil balance. Secondly, to allow trucking models to be run, so that the number of trucks required can be calculated.

The spoil balance issue will be further discussed in the solutions section of this chapter, but let's discuss trucking models now.

In my opinion, the time invested in running truck haulage models is another significant waste of engineering time. Trucking models simulate the haulage from excavation to dump location to determine typical cycle times and therefore the truck numbers required. An optimized version of a truck model will then theoretically smooth out the truck numbers required by building this consideration into the excavation scheduling. In other words the excavation schedule will be modified to smooth out the trucks required.

I believe there are alternate solutions for both of these two, very time consuming schedule additions and they will be discussed in the Solution section within this chapter.

Why Is Scheduling The Right Activities Critical?

There are two separate issues here. The first issue is a lack of inclusion of the right activities, which leads to a plan that is underdone. It misses critical issues, or doesn't tell people what they need to know, which leads to an ineffective plan and one that is doomed to failure when it is executed. Of the two issues, this is the more significant of the two. If your plan is going to be ineffective, some would say “why plan at all?”

The second issue is the inordinate amount of time wasted by engineers and time spent scheduling items that don't need to be scheduled. Over usage of engineering time results in an effective plan, but it is too effective, as it means wasting thousands of man hours. I've always believed in the ethos that engineering is about doing things smarter, not harder. In line with that, one of my favourite engineering themes is that it is not a case of whether the glass is half empty or half full, but that the glass is twice as big as it needs to be. Those sites regularly carrying out destination scheduling and truck haulage modelling, are working harder, not smarter and those engineering hours could be used so much more effectively elsewhere.

The Solution To Scheduling The Right Activities

Long Term Plans

The solution for long term plan, anything more than five years, is that you should be using an auto-scheduler if you have the software as an option and then scheduling using annual capacities. If you're not running an auto-scheduler, or you don't have access to one and therefore manually scheduling, then consider which equipment or processes you really need to schedule. Then schedule only those activities, and allow the schedule to generate the volumetric demand for the other activities. For example, in an open cut coal mine, you might schedule only the draglines and then allow this to determine how much drilling, blasting and truck/shovel excavation is required to meet the dragline schedule. I would describe this as scheduling by demand on the other activities.

Because the schedule driver will be ore production, in this process you are normally going to schedule either the mining of the ore itself, or the last step before the ore mining. So for example, the coal mining, or the prior activity, which is the dragline carrying out coal uncovering. But being a long term schedule, I wouldn't recommend scheduling an activity that has a large number of equipment items to schedule. Otherwise the schedule is going to take you a long time, so you need to strike a balance.

I don't believe that destination scheduling or haulage modelling are required in every long term plan, in fact I believe these two can be carried out using a different process. Firstly, every mine should have carried out some form of final landform modelling, although this is completely disconnected from any mine planning process. This landform modelling should be used to determine both the spoil fit and the optimum dump profile for final closure purposes. Finally, it can be used to determine which dumps each mining area should be hauled to so as to optimize the spoil fit and minimize haul distances.

The process is to carry out final landform modelling and resulting analysis at a life of mine level first, to determine the optimum location for all spoil on a whole of mine basis. Analysis is then carried out at a more detailed level by subdividing the

final landform model into bite sized chunks, which will vary in size, depending on the remaining life of the mine. But a typical example could be blocks that represent three years of excavation for each mining area. In this process, you can work out the average haul profile and volumes, and then weight average those areas to arrive at average cycle times during that period. From that process, it is possible to determine the trend in the cycle times over the life of mine schedule and feed this into financial models, along with the mining quantities, to establish the truck numbers required each year.

A final landform model such as this is a major exercise, so there is no need to repeat this work or recalculate the final landform model every year. I believe it sufficient to only review it in the following scenarios:

- upon completion of each of the bite sized chunk periods,
- if the final landform varies from reality by a long way,
- or something else has changed significantly.

In those situations there would be a need to carry out the final landform modelling again.

That of course is ignoring any license to mine requirements. Many governments demand information that mines can be held accountable against and therefore you may need to provide this information on a more regular basis than dictated by your mine planning. If you're forced to model landform more often for compliance purposes, that's just a bonus. But compliance is the driver, and not the mine planning or mine site accounting requirements.

Medium Term Plans

Once you are down at the level of mine plans of two years or less, then it's a different story. You are more likely to be scheduling all the equipment, as you need to, so that you have a higher degree of confidence in the plan outputs (which means you need to know the plan works). So you definitely need to be scheduling all activities, including drilling and blasting, but at a one to two year planning level, you can get by without scheduling some of the enablers, such a drill bench preparation. While you should be scheduling all equipment individually, I'm still a believer that you

don't need to destination schedule, or carry out trucking models, as that's already resolved in the final landform modelling process.

There will be plenty of people who will argue that as the Annual Plan feeds the annual budget, truck hours are critical to budget cost calculations, and therefore there is a need to run the truck modelling to get accurate hours. I would argue that this is just another example of where we continually overestimate the accuracy of our mine plans and place far too much reliance on these deterministic schedules. Typically mine plans are not that accurate and they don't reflect reality that well.

The previous methodology I have suggested for long term planning, basing it on final landform work and bite sized chunks will be fine. Given the variability and subsequent inaccuracy in our mine plans, I would suggest a better approach than truck haulage modelling would be reviewing how many trucks the mine site has currently and whether there will be a significant change in the haulage profile within the next year. If not, then the mine site is going to require a similar number of trucks for the budget. To determine truck hours, you know they are going to work at capacity, because they never work at less than capacity. So really it is simply a matter of applying the time usage model hours to the number of trucks to determine the quantity of truck hours for budget costing purposes.

Short Term Plans

Once we get into shorter term plans, say less than a year, then things change again. Short term plans at most mine sites are four months or less, and can be divided into two types. What I call short term plans, which I would classify as anything between two weeks and four months, and then anything less than two weeks, I call execution plans, which are discussed in the next section.

Short term plans are the same as the annual budget and two year schedule in that they should definitely include scheduling of all equipment and processes. The objective of a short term plan is to set the mine up for efficient implementation of the execution plan. Given the inherent variability at mine sites, it is important that the execution plan has flexibility built in and it is within the short term plan where you have the ability to create that flexibility.

Truck modelling is required to the level of giving the Production team trucking options, for example multiple dumps for one or more of the dig faces. That allows the Production team to select dump options based on how many trucks they are operating that shift. But to think that the short term plan can predict in detail how many trucks are going to be required every shift is crazy, there is way too much variability for that. So I don't see any point in trying to carry out destination scheduling or truck modelling in short term plans.

Execution Plans

And last but not least, we have the execution plan. Within the execution plan, you need to be scheduling all activities and equipment, but not only that, this plan is where you should also be scheduling key enabling activities, such as drill bench preparation, drainage activities, or ramp access works. The execution plan is about checking that you have the resources to get done what needs to be done and that there are no clashes between tasks. So at this level, the plan needs to include the scheduling of more equipment than all those longer time frame plans.

I would recommend that at execution plan level, mines should even be scheduling critical ancillary equipment such as dozers. Scheduling of ancillary equipment in this plan, is not so much about task quantities and completion rates and therefore how long the task will take for that ancillary equipment, but more so, it is about the allocation of equipment to specific tasks. Meaning that all Supervisors know what equipment they have to work with and can plan their shifts accordingly. In association with this, I would also suggest that for this reason, the execution plan needs to be at least a week in length.

Critical information that is derived from the execution plan, that needs to be clearly communicated, is what happens when equipment **isn't** available, such as it is down for maintenance. Is it a case that this team has multiple items of that equipment and so they have to work with one less item during that period? Or does this team have very few of that equipment and so it has been previously agreed that if there is a critical activity in action, such as drill bench preparation, then that team can attain resources from another team?

This is contingency planning and cannot be built into the schedule beforehand, as you can't plan for uncertainty like this. There are so many possibilities, but it is a case of laying out agreements and rules, and putting them in place each time an execution plan is created. This is one of the reasons for planning meetings, so that these issues can be discussed, coordinated and agreed upon.

Weather delays are similar in that they can't be built into the execution plan. I have come across mine sites who specify certain days as weather delays in their execution plans, for example Tuesday next week. This is to ensure that the execution plan aligns with their longer term plans, the annual budget might have included 4 days of weather delay in that month and so they include one day in each week. This is nonsense as weather days can't be scheduled unless you have a very reliable weather forecast for the next couple of days.

The whole concept of having to align different time frame plans on their time usage models and productivity assumptions is another mine planning fallacy. Longer term plans are "averaged" and should represent long term averages, whereas execution plans should represent the best forecast as to what is going to happen in the next week with the increased depth of knowledge that comes with planning a shorter time frame. For example, if an excavator is working on a wide face and in a well fragmented blast, then above average dig rates would be used.

12

CRIME 9: Theory of Constraints

In 2013, I came across two very smart gentlemen Kevin and Peter who were mining engineering consultants working for themselves. They spent some time telling me about the benefits of this new process that I hadn't come across before, it was called Theory of Constraints (TOC). Kevin and Peter explained TOC to me in a fair bit of detail and I could see that they were really passionate about it and that they really thought it could make a difference in the mining industry. I must admit, I thought it sounded like it had a lot of potential, but at the time I was on a mission. I was running my own consulting company Echelon Mining and had about 30 team members working for me and that was keeping me very busy. I was not looking to change direction for my business and to do something different, although TOC sounded interesting, it was just a distraction for me at that point in time.

But following that discussion, I always remembered TOC and always intended to look into it and find out a bit more about it one day. But the years went by and it never happened, I was busy in my own little area of the mining world and also, I don't think they had much success in trying to convince people to implement TOC. Unfortunately, it's a simple fact that the mining industry does not like change and Theory of Constraints was just too big of a change.

It was about 2018 before I recall starting to hear a few mining engineering colleagues that I greatly respected starting to mention TOC and they spoke about it in terms of the missed opportunity because it was not being incorporated. That reminded me and provided the prompt I needed to find out more about TOC. So I purchased the book titled "The Goal" by Eliyahu Goldratt and read it. Even though the story is about a factory production line, I could instantly see a place for TOC within the mining industry. I was sold.

What Is Theory of Constraints?

It would be remiss of me to write a book about mine planning and not include some discussion on TOC, at least in my mind anyway, as I think that it could be a game changer for the mining industry. There is a widespread lack of knowledge about TOC and subsequently I believe it is one of the most valuable and underutilised concepts that should be applied in mining. The awareness of TOC is slowly

growing, but it is struggling to gain any traction, I think that is because it is commonly misunderstood.

Mining, like many other industries, involves sequential activities which must be carried out so as to reach an end output from the system (ore shipped to customers), which, in TOC, is known as the system throughput. So mining is not dissimilar to a factory or production line, it is just likely that it involves significantly greater variability than many other processes.

I remember back in the 1990's when there was a big push in the mining industry for the introduction of Just in Time (JiT) processes, which came out of the lean manufacturing methodology. JiT involves the minimisation of inventories and batch processing, so products are produced "just in time". It became a "buzz word" in mining and there was much discussion about how the industry should operate on this basis. At the time, I remember thinking that it sounded entirely logical to apply it within mining, but it never really took off.

In hindsight, I believe there are two primary reasons JiT didn't become an accepted part of mining:

- The first is that the variability in the mining industry is so high, it is much greater than manufacturing, where JiT originated from. The high variability in combination with the balanced system that mines target, made it ineffective to run with the small inventories that are part of the JiT process.
- The second reason is that, in all honesty, in the mining industry, we have a general lack of understanding of inventories (or buffers in TOC speak), what their purpose is, and how we should use them.

Theory of Constraints was formulated by Eliyahu Goldratt in the 1980's and is a process methodology for maximising the throughput of manufacturing plants and other processes that involve a series of activities. TOC involves determining the system constraint, or what is commonly called the bottleneck, and focussing improvement efforts to better utilise that constraint so as to maximise system throughput. TOC involves five focussing steps in the methodology, as follows (these are further explained later in the chapter):

1. Identify the system's constraint;
2. Exploit the constraint;

3. Subordinate everything to the constraint;
4. Elevate the constraint; and,
5. Prevent inertia from becoming the constraint.

Effectively, it is a process of identifying the bottleneck, maximising the output of that bottleneck via operational process improvements, increasing bottleneck output through capital investment and then, upon production increase, check if it is still the bottleneck and repeat the process.

What Is The Issue?

Perhaps you've never heard of theory of constraints, or maybe you've heard of it, but only know a little about it. Or maybe, most frustratingly of all, you've put TOC forward as something that should be considered within your mining company, but you got nowhere with that hair brain scheme. If any of the above are true, then your mine is most likely operating in what is described within TOC as the cost world.

You're focussing on minimising cost by maximising the utilisation of resources – all of them. That's because if costs are at a fixed level and you can increase the throughput for that activity, then you reduce the unit cost of that activity and, if you can do that for every activity, then the overall unit cost must go down. But that's focussing on each individual activity, rather than the system as a whole and it's never that simple.

I have been involved in many improvement projects at mine sites where it is recognised that there is a bottleneck that is limiting production. Subsequently, the focus of the project has been to remove that bottleneck so that production can be increased, so it becomes all about removing bottlenecks. The problem with the standard approach taken by the mining industry is that removing the bottleneck and focussing on maximum utilisation leads to a balanced system and, in reality, a balanced system is disastrous.

Mines don't understand the value of a bottleneck, it is seen as a negative, a constraint that is holding the operation back. If you don't know how to use a bottleneck to control the remainder of the system and the throughput of the system as a whole, then the objective will always be to remove bottlenecks and you will subsequently

waste the opportunity that they present. In addition, mines generally don't really have an understanding of buffers (or inventories as we call them in the mining industry). Inventories are normally viewed as a necessary evil that are required to ensure that equipment can always work and, therefore, equipment utilisation can be maximised.

Consequently, the plan on which buffers are being managed is generally one of two approaches.

1. Attempting to minimise inventories as a whole, because they come at a high cost and involve money being spent earlier than it needs to. So the objective becomes to minimise inventories everywhere as much as possible, without having a negative impact on production.
2. Or the primary objective is to maximise utilisation; therefore, the site operates with high inventories for all activities. Management at these mine sites understand that to keep equipment working, and given the variability that exists in mining, then inventories are a necessary evil. So the mine site might run with huge inventories overall, but management believe that sacrifice is worth the value.

Personnel at the mine site do not necessarily think of the entire mine as a system and the system throughput as a whole, but instead focus on maximising the throughput for each of the individual activities.

Why Is Lack Of Use Of TOC An Issue?

The answer to that question is a very simple one: it is costing money.

As I'm sure most readers are aware, if costs are constant, then the greater the system throughput, the lower the unit cost of production. If you are not using a TOC approach, then you are currently sub-optimising and your system throughput is less than it could be. Those not using the TOC approach will tend to look at each individual activity and how it can be improved in isolation. TOC refers to this as local optima. Individual activities are improved based on the mistaken belief that the system will be improved, where this is not necessarily the case.

In addition, each activity is improved step by step without any consideration as to whether it is the bottleneck or not. You may be wasting both capital and labour on an activity that is not currently the constraint. Elevating an activity that is not the constraint will not increase system throughput and is therefore a waste of money. TOC focusses on system throughput rather than individual activity throughput.

The mining industry typically operates in a cost world and here is an example illustrating that. Let's take the example of a mine site with annual drill fleet capacity of 48 Mt and annual excavation capacity of 42 Mt. The mine site already has sufficient drilled and blasted inventory, but there is a mining area available for drilling that was planned to be drilled next month. Do you think the current monthly plan will schedule the drill fleet to park up, or to drill the area that was planned for next month? I would put my money on drilling the area that is available, because no one in the Production team wants equipment sitting around idle, it is in conflict with their KPI's and production targets. That is operating in a cost world.

Ironically, by focusing on maximising utilisation, mine sites are spending money earlier than they have to, therefore increasing costs. Even more so, the focus on maximising utilisation is primarily a focus on maximising the utilisation of mining equipment. Mining equipment represents capital that's already been spent (or sunk capital), so why is utilisation of sunk capital such an important driver for mine sites?

What Is The TOC Solution?

“Knowledge is Power” is a quote attributed to Sir Francis Bacon. You can't change what you don't understand, so before you can implement TOC an understanding of it is required. If you don't already have an understanding of TOC, I suggest you start with reading “The Goal” as previously referenced.

The first mine site I worked at when I started in the mining industry was Blair Athol Coal mine. Blair Athol was a one dragline mine, with that dragline responsible for the uncovering of all coal and there was no truck and shovel fleet required. Ironically, this mine operated on a TOC basis as there was a very obvious bottleneck, which was the dragline. This was widely understood and everyone on site clearly knew that the dragline was the one critical equipment at the mine and all priorities

were set around maximising its productivity. So quite by accident, I started mining on a Theory of Constraints basis thirty years ago.

Blair Athol is a simple example, however, and the process is far more complex if you have a large mine with multiple operating pits and excavation units. In that situation, don't position the excavation fleet, or anything else with multiple equipment items (such as drills), as the constraint. It is too challenging to run an activity with multiple equipment items as the constraint due to the change management process required. In addition, the possibility of the constraint being managed by multiple supervision teams on site further complicates the process. The best option is to position fixed plant or infrastructure as the constraint, for example, the processing plant, railing system or the port. In each of those cases, it is just one item (or system) that requires management.

The first step in TOC is to identify the constraint, so let's use an example to highlight this process and how we would apply it. In this example, we have an open cut coal mine with the following production capacities.

1. Drilling fleet annual capacity of 1,450 km;
2. Blasting contract for annual supply of 63,000 tonnes of explosives product;
3. Annual truck and shovel fleet capacity is 37 Mbcm;
4. Annual dragline capacity is 90 Mbcm;
5. Coal mining fleet annual capacity is 9 Mt;
6. Coal processing plant annual feed capacity is 8.5 Mt;
7. The annual railing contract is for 8.9 Mt; and,
8. The annual port contract is for 9.6 Mt.

This is the fundamental information required to identify the constraint; however, further data is required and that additional data is included in Table 12 - 1.

With all of that data, it is then possible to carry out the process of identifying the constraint. I've completed the process, along with some explanation and included this in Table 12 - 2. In this example, the dragline fleet is the constraint, closely followed by the drilling fleet. We have now identified the constraint, but what's just as important is we also understand the level of excess capacity in all the other activities on site. We know those activities which are close to being the constraint and, therefore, under certain circumstances could become the constraint.

Parameter	Unit	Value
Average Pattern Area	m ²	63
Average Powder Factor	kg/bcm	0.6
Strip Ratio	bcm/t	11.5
Average Waste Depth	m	80
Average DRE Depth	m	50
Dragline Rehandle	%	0.7
Yield	%	0.87

Table 12 - 1 Production Parameters

This is important because the term balanced system refers to those systems without a specific constraint and where numerous or all activities within the system have equal capacity. Because the mining industry generally operates within a cost world, it aims to maximise the utilisation of all equipment and resources and therefore aims to install balanced systems, where everything in the system is always operating at maximum utilisation. This sounds good in theory, but is detrimental as the constraint regularly switches between activities, creating an unstable system.

Table 12 - 2 shows that the drilling fleet has 8% excess capacity, the CHPP has 16% excess capacity and the truck and shovel fleet has 17% excess capacity. Given the level of variability within mine sites, there will definitely be periods where drilling, CHPP or the truck and shovel fleet could be the constraint. While this example doesn't represent a balanced system, it is evident how easily the constraint would move between activities. Once you've carried out this exercise, you will then have an understanding of how close your system is to a balanced system.

Activity	Description	Product Tonnage Capacity (Mtpa)	Excess Capacity (%)
Port	Dictated by Port contract of 9.6Mtpa	9.6	50%
Rail	Dictated by Rail contract of 8.9Mtpa	8.9	39%
CHPP	CHPP feed capacity = 8.5Mtpa	7.4	16%
	Yield = 87%		
	Therefore CHPP product capacity = $8.5 \times 0.87 = 7.4$ Mtpa		
Coal Mining	Coal mining fleet capacity = 9Mtpa	7.8	22%
	Yield = 87%		
	Therefore CHPP product capacity = $9 \times 0.87 = 7.8$ Mtpa		
Dragline	Dragline capacity = 90 Mbcmpa	6.4	0%
	Dragline Rehandle = 70%		
	Therefore Dragline prime = $90 / (100\% + 70\%) = 53$ Mbcmpa		
	Dragline prime = 50m of 80m total waste depth		
	Therefore dragline prime = $50 / 80 = 62.5\%$ of waste		
	Total waste to match dragline rate = $53 / 0.625 = 85$ Mbmpa		
	Strip ratio = 11.5:1 therefore		
	Dragline annual coal uncover = $85 / 11.5 = 7.4$ Mtpa		
	Yield = 87%		
	Therefore product coal capacity = $7.4 \times 0.87 = 6.4$ Mtpa		
Truck & Shovel	Truck/shovel waste = prime less dragline waste = $100 - 62.5 = 37.5\%$	7.5	17%
	Truck/shovel capacity = 37 Mbcmpa		
	Therefore equivalent total prime capacity = $37 / 37.5\% = 99$ Mbcmpa		
	Strip ratio = 11.5:1 therefore		
	Truck/shovel annual coal uncover = $99 / 11.5 = 8.6$ Mtpa		
	Yield = 87%		
	Therefore product coal capacity = $8.6 \times 0.87 = 7.5$ Mtpa		
Blasting	Annual blast contract = 63 kt	7.9	23%
	Average Powder Factor = 0.6kg/bcm		
	Therefore annual waste blasted = $63 / 0.6 = 105$ Mbcmpa		
	Strip ratio = 11.5:1 therefore		
	Coal uncovered by blasted waste = $105 / 11.5 = 9.1$ Mtpa		
	Yield = 87%		
Therefore product coal capacity = $9.1 \times 0.87 = 7.9$ Mtpa			
Drilling	Drilling annual capacity = 1,450 km	6.9	8%
	Pattern area = $63m^2$		
	Therefore waste drilled capacity = $1,450 \times 63 / 1000 = 91$ Mbcmpa		
	Strip ratio = 11.5:1 therefore		
	Coal uncovered by drilled waste = $91 / 11.5 = 7.9$ Mtpa		
	Yield = 87%		
	Therefore product coal capacity = $7.9 \times 0.87 = 6.9$ Mtpa		

Table 12 - 2 Identifying The Constraint

You also have an understanding of two key areas

1. Where improvement efforts should be focused; and,
2. The critical activities in your mine schedule.

This book is not about continuous improvement in production, it's about mine planning improvement, so we'll limit the amount of discussion on production continuous improvement. It is obvious from Table 12 - 2 that continuous improvement applied to a process that is not the bottleneck, for example the port or railing, will not increase overall throughput. I'm not saying that improvement efforts shouldn't be applied to non-constraints, but understand that it won't increase the system throughput and it will be more beneficial to follow the TOC process instead.

So let's understand the five TOC focussing steps in more detail from a mine planning perspective.

1. Identify the constraint. We just completed this exercise in Table 12 - 2
2. Exploit the constraint. Obtain as much capacity from the constraint as possible without substantial capital investment. Production depends on two elements, operating hours and productivity, this step involves analysing how both can be increased. There are a number of aspects of both elements that are impacted by mine planning. For example, mine design can maximise productivity by designing appropriate width and height benches, scheduling can minimise operating delays via shorter or less frequent deadheading, or maintenance may be scheduled on constraint equipment so that it coincides with other operating delays.
3. Subordinate everything else to the constraint. Make the constraint (in this case the dragline fleet) the number one activity on site and subordinate everything else to it. If there are any schedule conflicts when scheduling, then the constraint activity takes priority. But, in fact, if you're scheduling knowing the draglines are priority, then you'll actually schedule so that they are not impacted by any other activities. The maintenance and ore processing planners will also plan accordingly.
4. Elevate the constraint. Now investigate how the capacity of the constraint can be further lifted. Typically, this is usually via capital investment, such as larger dragline buckets or upgrades to electrical systems.

5. Start again. In elevating the initial constraint (the dragline), production may increase sufficiently that the constraint moves to another activity. In the example used, if dragline production increased by greater than 8%, then the drilling fleet now becomes the constraint. So reassess the constraints and start the process at Step 2 again.

TOC provides a map for the application of continuous improvement once you know the constraint. But what is particularly important for mine planners is that we can see how scheduling plays a key role in implementing TOC. Understanding TOC and knowing what the constraint is, facilitates the mine planner taking a different approach to mine planning from the standard mindset. There is the potential for substantial upside in the production capability of the constraint activity if all parties on site (planning, production, maintenance, etc) work together in prioritising this equipment.

This is just the start and I've barely scratched the surface of TOC in this chapter, primarily because I am an amateur at TOC with no formal training in it – but I am a believer. There's much more to it and so again, I'd recommend reading “The Goal” and then seeking further education in this area (or employ someone who is knowledgeable in TOC). It is particularly relevant if designing a new mine from scratch. I would analyse two options:

- The standard mining approach - putting a balanced system in place where all equipment is sized to maximise utilisation and be no larger than required for a balanced system. Design inventories as they would typically run within a standard mine site.
- The TOC option – design the mine site with a targeted constraint, selected on the basis of which activity is most appropriate as the constraint. Design in buffers (inventories) as they would be utilised within the TOC approach.

Since I've begun discussing the benefits of TOC in some of my LinkedIn articles, I've had many discussions with Hendrik Lourens from Stratflow about the application of TOC in existing mines. Stratflow have implemented TOC in more than 85 mining interventions over the last 20 years, with throughput and cost improvements of typically 20% or higher. The increased stability in production also improves the stability of the mine plan and frees up time and resources. Experienced miners know that this is a prerequisite for enhancing safety.

The positive effect that a focus on global performance versus local optimisation has on people and culture is often missed. Experience shows that the success achieved by aligning the activities of departments towards a common goal often results in better coordination and employees becoming more proactive. Employees gain control of the outcome in their area of responsibility and subsequently engagement increases.

Having a basic understanding of TOC, and with over thirty years of experience observing how mine sites work, I am confident those levels of improvement are achievable and that TOC definitely has a place in the mining industry.

13

**CRIME 10: Mine Plans Are
NOT A Performance Driver!**

While I was working in mine planning at one of my mine sites we had a new General Manager take over management. He was extremely ambitious and driven and he definitely managed by using numbers, it was all about the numbers. He had ambitious targets for the mine and, for him, the mine plan was the way to drive performance and achieve those targets. His approach was to set ambitious targets for equipment production rates, so if, for example, a shovel typically targeted 30,000bcm per day, he wanted to set the target at least 30% greater than that at 39,000bcm per day.

To exacerbate the issue even further, he wanted to call those daily targets “the plan” because he actually wanted people to just view that as the norm. Of course, this caused many issues and a lot of discussion about the fact that “plan” and “mine plan” were already commonly used terms throughout the industry already and at this mine site, and we couldn't start calling something that was 30% above normal “the plan”. I pushed back hard on this issue and eventually he agreed to the use a different term instead, but the best I could do was to get him to use the term “schedule” instead of “plan”, so the schedule targets became extremely ambitious.

However, I was successful in managing to achieve a disconnect between what was described at this mine as the “schedule targets” and the “assumed production rates” I used in all of the mine plans. The disconnect was on the basis that the “schedule targets”, as per the General Manager’s requirements, were only ever measured on a daily basis. Every day was compared against the “schedule target”, so they effectively became a benchmarking process and these “schedule targets” were never incorporated into any of the planning. As I’ll discuss further in this Chapter, this is an extremely important disconnect.

What Is The Problem?

The mining industry, like many other industries, focusses heavily on continuous improvement as part of standard operating practices. To assist in embedding continuous improvement into operations, it is common to incorporate higher target production rates, ore recovery rates and other improvements into mine plans. Effectively, the mine plan becomes a tool used in the process of driving productivity

improvements. Management don't want to put out a mine plan with input assumptions that only match what actually happens on site. The belief is that this would drive a culture to “just maintain the status quo” and this would become acceptable when clearly it is not. They want to drive improvement by using higher production rates in the mine plans, so that these become the “norm”.

Here’s a perfect example of this. I carried out consulting work once for a very large mining organisation who had a process that required that a 2% production improvement every year was built into their long term mine plans. There were no calculations or science behind the 2% assumption used, it was purely an acknowledgement of the fact that mines have historically improved their production rates and costs so this should continue into the future. Who was I to point out that much of this historical increase in production rates was due to the trend to using larger equipment so as to improve labour productivity? And that the rate of increase in equipment capacity growth was slowing, as it should do.

A 2% production increase every year may sound reasonable, that is until you actually look at the numbers. As an example, a dragline historically producing 12 Mbcm per annum, if inflated at 2% per annum, would be planned at the production rates as shown in Table 13 - 1, highlighting the growth over time.

When	Production
Now	12.0
In 5 Years	13.2
In 10 Years	14.6
In 20 Years	17.8
In 30 Years	21.7
In 50 Years	32.3

Table 13 - 1 Impact Of Continuous Improvement

As you can see, the increases are substantial and do get to the point of being absolutely absurd. It may be conceivable that this particular sized dragline could increase production from 12 Mbcm to 14.6 Mbcm per annum over a 10 year period, but then that is likely around the maximum possible potential for this size machine. Anything above that in production is likely going to need a larger sized dragline at a substantial capital investment (which of course was not in the capital budget). The mine had a life of greater than 50 years so the plan used values of +30 Mbcm in the later mine life. I'm not sure that there is a dragline anywhere in the world moving 30 Mbcm of material per year.

This example highlights an issue I have seen regularly within the mining industry, that is one of building improvement assumptions into the mine plan or budget, without an actual plan as to how those improvements are going to be achieved. There is a philosophy of, "we need to achieve these improvements to survive so, somehow, we'll just make it happen". But unfortunately, the industry has a long history of under-achieving on those targeted improvements.

A related issue I have come across regularly at mining companies is one of "market expectations". This goes along the lines of, "we're not going to provide guidance to the market this year that is worse than last year, so re-work your plan and budget and provide us with a better set of results." I've seen this commonly happen once the mining engineers have completed the first draft of the mine plan and it is forwarded to executive management for approval.

The worst example of this I have seen involved the first level of management review requesting improved numbers to forward to higher management, so by the time it had been approved by lower level management, it was already a plan that the planning engineers didn't believe was achievable. But worse still, before it was sent to top level management for review, a corporate team enhanced the results further by applying improvement overlays within spreadsheets so, at that stage, there was effectively zero probability the plan would be achieved.

Being an engineer myself, I believe that engineers graduate from university programmed to apply a conservative approach to design, planning and budgeting. Maybe this is more applicable in some disciplines than others, for example, I studied Civil Engineering and the factors of safety are very high because if something goes

wrong and the structure fails, people die as a result. But overall, I think engineers are programmed to operate in a conservative manner.

So I have been surprised when I've seen the level of optimism built into some mine plans by the planners. This has not been around the productivity assumptions, but more so around the gaps left between activities but other assumptions. For example, the time between when a dragline blast is fired and the blast is dozed and prepared for the dragline to walk onto the blast and start excavation. I think this largely stems from the conservatism being “whipped” out of them following years of being told to produce plans with better results.

Why Are Optimistic Plans A Problem?

Let's not beat around the bush here, optimistic mine plans are an **insidious disease** that causes great harm to the industry. Optimistic planning assumptions, in my opinion, are one of the two biggest issues in mine planning (along with deterministic scheduling) and are one of the primary reasons that mine plans are viewed as a joke and not worth the paper they are printed on.

The problem with optimistic assumptions are that they create unrealistic schedules, but worse than that, they create optimistic schedules and these are the worst kind of schedule. Because they overstate the outcomes (primarily the amount of ore that will be produced), these outcomes then become an expectation from everyone on site. So rather than being a plan for what might happen, it just becomes a vision of what management would like to happen.

I remember during my time with Thiess Contractors, the company had a very strong culture of “it's always better to surprise on the upside”. And that was a critical driver behind all communications, including mine plans (which are just a form of communication). But if your mine plan includes optimistic assumptions, then you have locked yourself into surprising on the downside instead and that is never a pleasant experience.

I once carried out an audit of the Technical Services Department at a very large Australian coal mine. I spent the first couple of days asking questions of everyone in Tech Services, along with a number of their customers (Production, Coal

Processing, Maintenance, etc). One very common theme that came out of my questioning was how unrealistic everyone thought the mine plans were. There was a general annoyance at the fact that they were unachievable and it was routine practice at this mine that at some point it was inevitable that shortcuts would have to be taken to try and achieve the plan. Labour resources were being wasted in scrambling, re-designing and re-scheduling to try and meet an optimistic plan, rather than invested in producing a smarter plan in the first place.

This example is not uncommon in Australian mining and it creates a cycle as follows, that I have experienced at numerous mine sites:

- Mine plan produced with both planning and operations personnel knowing it is not achievable.
- Operations follow the plan initially, slowly getting behind plan, but that hasn't created a problem yet.
- A trigger causes a change to the plan, that trigger typically being one of two issues:
 - A clash between two planned tasks, so one task can't commence because a task it is dependent on is not complete yet.
 - Or, it becomes clearly inevitable that continuing to follow the plan is not going to achieve the plan targets, which is generally represented by mining sufficient ore.
- The plan is modified or a new plan is produced. That plan will involve one or more of the following changes, which I will group under the terminology "shortcuts".
 - One or more items of equipment will be relocated from their current mining area to a new area, resulting in time wasted relocating.
 - Designs will be modified, leading to equipment being operated in a lower productivity manner. For example, a blast will be reduced in width to the key area only, leading to a higher percentage of drill walk time on the blast, or lower excavator productivity due to a reduced face width.
 - A non-urgent task (or tasks), will be sacrificed for a more urgent task, one that directly leads to plan targets being achieved.

- Ore will be blended in a sub-optimal manner due to insufficient inventory available, or the need to increase product tonnages leads to inferior products being processed at higher recoveries.
- The plan modifications result in the plan being achieved, or very close to being achieved, close enough that it doesn't present any issues.
- When scheduling the next plan, it is apparent that it is more difficult to achieve targets. This is because the plan is starting on the back foot, for example:
 - inventories have been drawn down,
 - the non-productive material that was not scheduled in the re-sequence of the previous plan is now in the way and has to be sequenced at the start of this plan,
 - non-critical tasks have been delayed, meaning that these tasks now become critical and add to the number of critical tasks in the plan (increasing risk),
 - in total, reduced quantities have been moved compared to plan, due to time wasted relocating or operating in lower productivity modes, so by default the mine is starting behind plan.
- The cycle continues of having to take shortcuts or make other modifications to the plan so as to achieve schedule targets, but each time those modifications scale up in intensity. It continues until it is no longer possible to achieve targets via shortcuts, at which time the mine site finally admits defeat and there is a major write down in the plan and budget.

Sound familiar? The problem with optimistic schedules is that they ultimately force the mine site to take shortcuts and the shortcut is never a better solution. If the shortcut was the best solution, it would have been built into the original plan in the first place. Shortcuts generally result in a reduction in efficiency and are just a method of deferring the pain (involved with not hitting target). But while deferring the pain, they also increase the ultimate size of the issue at the same time. Instead of frequent small surprises on the downside, the process leads to a fewer number of large surprises on the downside. But the irony is that shortcuts ultimately lead to reduced production over the long term as they reduce the mine site efficiency.

I find the best way to reinforce a point is through numbers and an example is perfect for that. Let's expand on the discussion from Chapter 5 on slippage in the schedule,

where we have two tasks in series in the schedule (one must be finished before the other can start), the example used in that chapter involved blasting and excavation. In the scenario that our assumptions match historical performance, then we are effectively using P50 assumptions (50% probability of finishing on or ahead of time). As shown in Chapter 5, P50 assumptions for each of the two tasks in series lead to only a 25% probability of finishing either ahead of plan, or on plan.

But were we using P40 assumptions instead, then the overall probability of finishing ahead of plan reduces to 16% or, alternatively, if we were to use P60 assumptions, then the probability of finishing ahead of plan increases to 36%.

Table 13 - 2 displays the change in probability of finishing on time or ahead of plan as the input assumption probability changes. The impact is significant in terms of reduction in the probability of finishing ahead of plan. For example, halving the input probability from P60 to P30 leads to a four times reduction in the probability of finishing ahead of plan – from 36% to 9%.

Input Probability	"Ahead of Plan" Probability
P90	81%
P80	64%
P70	49%
P60	36%
P50	25%
P40	16%
P30	9%
P20	4%
P10	1%

Table 13 - 2 Probability of Finishing Ahead of Plan – Two Tasks

Note that the simple example above doesn't consider two very important parameters.

1. Open cut mining processes normally involve about six tasks in series rather than two as used in the example (drill, blast, excavate waste, mine ore, process ore, rail ore). This significantly reduces the probability of all tasks

finishing on or ahead of time. A P50 scenario becomes 0.56 which equals 1.6% probability, rather than 0.52 for two tasks which is the 25% displayed in Table 13 - 2. The resulting probabilities for six consecutive tasks are shown in Table 13 - 3, which highlights the significant impact of using lower probability input assumptions. Halving the input probability from P60 to P30 now leads to a sixty four times reduction in the probability of finishing ahead of plan, from 5% to 0.1%. That is material!

2. All of the above analysis excludes consideration of the impact of inventories, which are used for the purposes of introducing buffers between tasks. Those buffers serve to unlink tasks that are in series, so that one task starting is not dependent on the preceding task being completed. The introduction of buffers will significantly increase the probabilities of finishing ahead of plan from those numbers shown in Table 13 - 3. These are complex calculations and not something I'm going to cover further in this book. But it is a topic I am passionate about so watch my website or LinkedIn for further blogs on this subject. However, the one conclusion we can draw from Table 13 - 3 is that planning with higher levels of optimism in the input assumptions requires increased inventory levels to mitigate the increased risk and uncertainty (lower output probabilities).

Input Probability	"Ahead of Plan" Probability (2 Tasks)	"Ahead of Plan" Probability (6 Tasks)
P90	81%	53%
P80	64%	26%
P70	49%	12%
P60	36%	5%
P50	25%	2%
P40	16%	0.4%
P30	9%	0.1%
P20	4%	0.01%
P10	1%	0.0001%

Table 13 - 3 Probability of Finishing Ahead of Plan – Six Tasks

Why else is optimistic planning an issue? As discussed a number of times throughout this book, mine planning at many mines today involves additional processes such as optimisation, destination scheduling and haulage modelling. Often mine plans require many hours of labour to complete all of the processes involved. I cannot think of anything that is a larger waste of time than carrying out optimisation on a plan that is optimistic and the planners know is never going to be achieved. It is a similar situation with destination scheduling and haulage modelling. How demoralising is it to put all that time and effort into a plan you know is a waste of time? And what a giant waste of precious resources, that engineering time could be much better invested in other ways to improve the operation.

Mine plans typically have a short shelf life already due to the large range of variability that occurs within mining processes, we don't need to exacerbate this issue and shorten shelf lives even further by using optimistic assumptions. That is exactly what these assumptions do, they shorten the shelf life of the plan. Optimistic plans help to drive the negative culture that mine planning is a waste of time.

The Solution To Optimistic Planning

There are a range of solutions to this issue, but like most change, it starts with a change of mind set and, subsequently, a change in culture. The first step required is a separation of mining performance drivers from mine planning. This is best done by distinctly separating the two into what I would label as benchmarking and mine planning. Now I'm not saying that targets shouldn't be set and that there shouldn't be a focus on continuous improvement and that numbers are used as part of that process. But what I am saying is that this is not the job of the mine plan, this is the job of benchmarking.

So sure, send out daily performance results to stakeholders and have charts up on the wall at the pre-start area or in meeting rooms. But don't show actuals compared against an optimistic mine plan or budget, instead of labelling the target line as "Plan" or "Budget", call it exactly what it is – "Target", or "Benchmark" could be another term. Then have another set of results or charts that reflect how the mine is actually going against Plan.

I would like to point out an area of caution to be applied here though. Setting high benchmark targets and potentially rewarding the relevant personnel for high performance leads to the distinct possibility of increasing the variability of operations at your mine site. I can't tell you the number of times I have seen excavation equipment achieve record production during one shift with all the ceremony and acknowledgement that come with it. But then no one ever talks about the following shift, where faces were left in a mess or the "cream" dirt had been dug and the unproductive material was left behind. They never talk about the production on the next shift and how far below plan it was.

I've seen mine sites distribute stickers or belt buckles celebrating a record shift or day of production from a dragline or shovel where it shows the tonnes moved and the date. But I've never seen a mine site celebrate consistency of production, which is what I'd rather see. Consistency of production lowers mine site variability, therefore making plans more reliable and able to be followed and reducing inventory requirements and the need for flexibility (additional mining faces). What about a sticker instead that celebrates "seven days in a row with a standard deviation less than 5,000 bcm". Yeah I know, that was tongue in cheek, no one understands or gets excited about standard deviations. Maybe "fifteen days in a row at or above plan"? That's consistency.

I believe there is a strong need to create a culture where the mine plan should be the most realistic estimate of what is actually going to happen, but that it is a base line and the mine site is always striving to perform above plan. This story from my time spent working for Thiess Contractors at Burton Coal Mine exemplifies this perfectly.

Moving from a large corporate miner, to working for a contractor was a refreshing change and the source of plenty of learning. One of the biggest differences that I noticed was a significant cultural difference between the two in the area of mine planning. Large corporate miners (and I still see this today at mining companies) would construct mine plans based on ambitious production parameters because they believed that using ambitious targets within the plan would drive performance behaviour change. This leads to optimistic mine plans that would cause site to spend all their time chasing their tail and taking shortcuts in trying to achieve those mine plan targets.

But the contractor had a totally different approach. The contractor, in submitting their tender for an operation, incorporated the most conservative set of assumptions and targets they believed they could get away with and still win the contract. Therefore, we had a plan with a set of targets that was not overly optimistic nor impossible to achieve, because if that was the scenario, then the contractor knew that we were going to lose money so wouldn't tender for the work.

With a conservative set of mine planning assumptions that we based our scheduling on, the mining engineering team got to spend much more of our time focusing on smarter planning and how we could do better than the plan. Whereas, within the corporate miner, where we had optimistic mine plans, we spent our time trying to work out how on earth we were going to meet that plan and what shortcuts we might need to take.

I implore you, for the sake of adding some credibility to mine planning, please separate benchmarking from mine planning.....

But that discussion is all site based and effectively about mine management using the plan to drive performance, there is another area where the planning process is broken and that is at the executive level within mining companies. The problem stems from market guidance and plans that have previously been communicated to shareholders, which have formed market expectations.

As discussed earlier in this chapter, there is a need to separate the two processes of mine planning and any adjustments made to satisfy market expectations. If adjustments to plan outputs are required for the purposes of improving guidance to the market, make this a process that happens post plan approval within the company and nothing to do with the plan that site is being measured against and held accountable to. Mining companies will never achieve buy in of the plan and ownership by site personnel charged with executing it, if there is a head office team who make adjustments to the outputs after the site has created the best plan they think they can achieve.

Reference Class Forecasting is a methodology that can be used to stop the ongoing process of optimistic mine planning. I had not heard of it until a colleague sent me a paper on the subject in 2020 and it is something that should be used more frequently in the mining industry. That paper is called "Curbing Optimism Bias and Strategic Misrepresentation in Planning: Reference Class Forecasting in Practice"

by Bent Flyvbjerg and was written in 2006. Reference class forecasting is based on theories of planning and decision making that won the 2002 Nobel Prize in economics, so we should sit up and take notice.

The paper is a fascinating read on the historical and repeated inaccuracy of both cost and time projections for public works projects in the United States. It is based on the “planning fallacy”, a systematic fallacy in planning and decision making under which people underestimate the costs, completion times and risks of planned actions. The recommended solution to this issue is to instead use actual performance in a “reference class” of comparable actions previously carried out, as the basis for the forecast.

Two previous studies had documented that forecasts of cost, demand and other impacts on major plans and projects have remained constantly and remarkably inaccurate for decades. For example, inaccuracy in cost forecasts for transportation infrastructure projects was found to be, on average, 44.7% for rail, 33.8% for bridges and tunnels and 20.4% for roads. Sound familiar?

There is a great discussion within the paper by Flyvbjerg on the issue often being bias, rather than inaccuracy and how, for a range of reasons, it is frequently an actual choice to bias the assumptions used. They describe this as strategic misrepresentation, which is just a nice way of saying “lying”!

This is somewhat akin to what I have come across many times in the mining industry. That is researching the mine’s historical performance to ascertain typical production times and rates. But then management direct that rates should be increased, without there being a clear plan as to how those increases will be achieved, just that the plan needs to show it. The terminology used to explain this typically is “market expectations”.

Reference class forecasting for a particular project requires the following three steps:

1. Identification of a relevant reference class of past, similar projects. The class must be broad enough to be statistically meaningful but narrow enough to be truly comparable with the specific project.

2. Establishing a probability distribution for the selected reference class. This requires access to credible, empirical data for a sufficient number of projects within the reference class to make statistically meaningful conclusions.
3. Comparing the specific project with the reference class distribution, in order to establish the most likely outcome for the specific project.

Thus, reference class forecasting does not try to forecast the specific uncertain events that will affect the particular project, but instead places the project in a statistical distribution of outcomes from the class of reference projects.

I do love a good story and the following story from Flyvbjerg's paper explains it very well.

Daniel Kahneman relates the following story about curriculum planning to illustrate how reference class forecasting works. Some years ago, Kahneman was involved in a project to develop a curriculum for a new subject area for high schools in Israel. The project was carried out by a team of academics and teachers. In time, the team began to discuss how long the project would take to complete.

Everyone on the team was asked to write on a slip of paper the number of months needed to finish and report the project, the estimates ranged from 18 to 30 months. One of the team members—a distinguished expert in curriculum development—was then posed a challenge by another team member to recall as many projects similar to theirs as possible. “How long did it take them at that point to reach completion?” the expert was asked. After a while he answered, with some discomfort, that not all the comparable teams he could think of completed their task. About 40% of them eventually gave up. Of those remaining, the expert could not think of any that completed their task in less than 7 years nor of any that took more than 10. The expert was then asked if he had reason to believe that the present team was more skilled in curriculum development than the earlier ones had been. The expert said no, he did not see any relevant factor that distinguished this team favourably from the teams he had been thinking about.

The wise decision at this point would probably have been for the team to break up, according to Kahneman. Instead, the members ignored the pessimistic information and proceeded with the project. They finally completed the project 8 years later, and their efforts went largely wasted—the resulting curriculum was rarely used.

In this example, the curriculum expert made two forecasts for the same problem and arrived at very different answers. The first forecast was the inside view; the second was the outside view (the reference class forecast). The inside view is the one that the expert and the other team members adopted. They made forecasts by focusing tightly on the project at hand, considering its objective, the resources they brought to it and the obstacles to its completion. They constructed in their minds scenarios of their coming progress and extrapolated current trends into the future. The resulting forecasts, even the most conservative ones, were overly optimistic.

The outside view is the one provoked by the question to the curriculum expert. It completely ignored the details of the project at hand and it involved no attempt at forecasting the events that would influence the project's future course. Instead, it examined the experiences of a class of similar projects, laid out a rough distribution of outcomes for this reference class, then positioned the current project in that distribution. The resulting forecast, as it turned out, was much more accurate.

Let's explain reference class forecasting via an example. If the public works were to estimate time and cost required to construct a bridge, this is normally built up from first principles. So all of the items will have basic productivity and time assumptions applied and then be combined together into a schedule. These times and rates will typically come either from theoretical calculations or from historical past performances.

What reference class forecasting does is review historical data for similar types of bridge construction and works out a distribution range for time and cost overruns. They would use this analysis to arrive at data similar to that shown in Figure 13 - 1, which is drawn directly from the paper. For the historical projects, 20% of them had an overrun of 0% or less, 60% of the projects had a maximum overrun of 20% and the entire set (100%) had a maximum overrun of 65%.

Based on these probability distributions, the required uplifts needed to carry out the third step in the reference class forecast may be calculated as shown in Figure 13 - 2. The lower the acceptable risk of overrun, the higher the uplift in forecast required. For example, if the acceptable risk of overrun required is 50%, then an uplift in estimate of 15% should be applied. However, if the acceptable risk of overrun is only 10%, then an uplift in estimate of 50% is required.

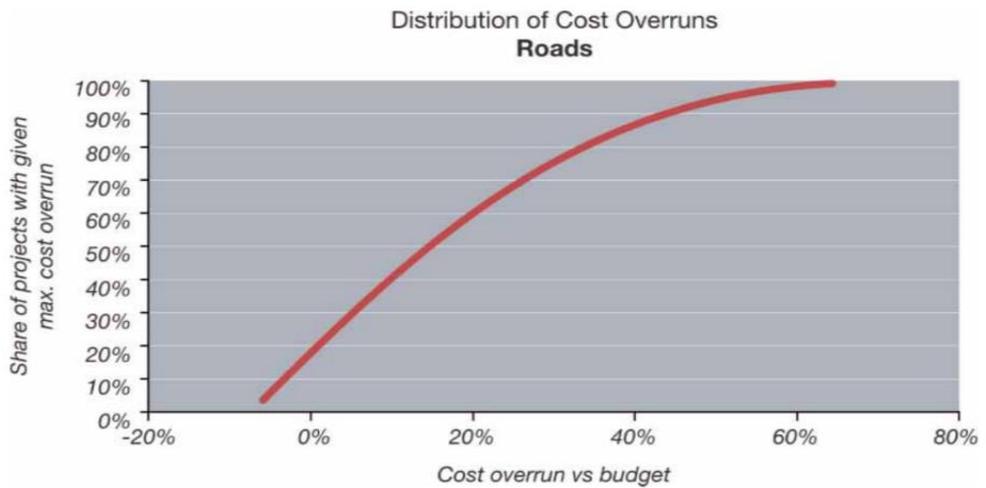


Figure 13 - 1 Probability Distribution of Cost Overrun

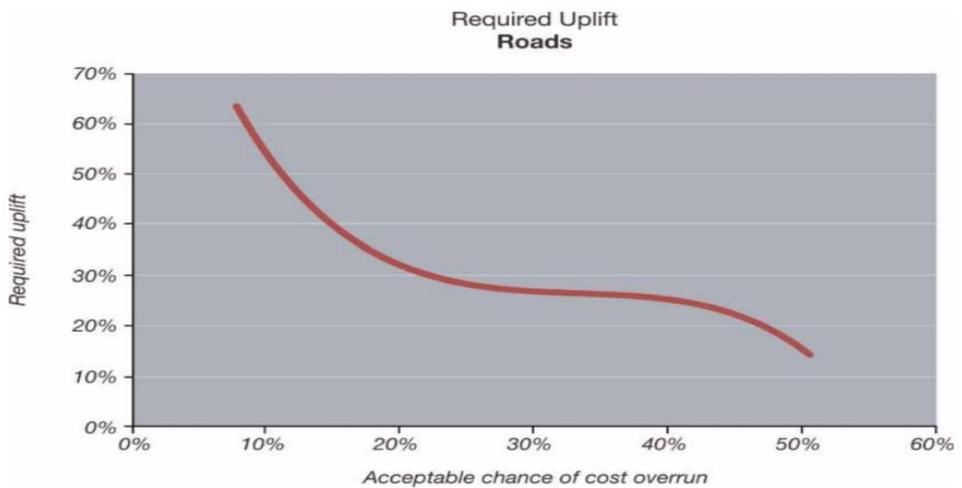


Figure 13 - 2 Required Uplift

Now my question is this, why can't we apply a reference class forecasting approach within the mining industry? In fact, quite by accident I carried out a process similar to this only about a month before reading this paper and I relate that example in the

following story. But had I already read this paper, I likely would have gone about that review with a slightly different approach.

I was asked by a mining company to review their recently completed Quarterly Plan, as they wanted to know with a high degree of confidence the quantity of ore they were going to have to blend with so they could maximise their blending efficiency. This was a complex mine with over twenty items of excavation equipment and seven different operating pits. I had not carried out any mine planning at this mine previously, so my lack of familiarity with the mine along with its complexity placed me at a difficult starting point. It meant I couldn't work through the schedule at the level of detail required to provide any commentary as to my level of confidence in their schedule; at least not without taking a large amount of time to get sufficiently familiar with the site and their mine plans.

This is not an uncommon situation as a consultant so you often have to “engineer” a solution. I reviewed the data I had available and found I had the following:

- The Annual Mine Plan & Budget for the mine.
- The Quarterly Plans were updated each month, so effectively a rolling three month plan. I was provided the twelve Quarterly plans that corresponded with the Annual Mine Plan period.
- Actual production results for each of the 12 months during the Annual Mine Plan period.

The question I had to answer was “what is your confidence we will mine the coal that is reported in our most recent Quarterly Plan?” It was widely recognised at the mine site that the Annual Plan was based on optimistic assumptions, but it was believed that the Quarterly Plan was more realistic as it was based on historical operating hours and productivities. In fact, because of the use of historicals, the mine site itself had high confidence in their Quarterly Plans.

The first step I carried out was to compare the Annual Plan and Quarterly Plans by summing the first month of each of the twelve Quarterly Plans and then comparing those totals against the Annual Plan. This confirmed the optimism built into the Annual Plan, as summing the twelve Quarterly Plans led to a dragline total volume that was 90% of the Annual Plan, a truck and shovel total that was 88%, but surprisingly a coal mined total that was 99% of Annual Plan.

The next step was to compare the totals of the twelve Quarterly Plans against the actual quantities over that twelve month period. The results were that actual performance for dragline prime was 96% of the Quarterly Plan total, truck and shovel excavation was 94% of plan and coal mined was 91% of plan. There were a number of conclusions possible from that analysis, but the primary two observations were as follows:

1. The first stems from the fact that historical production rates were used in the plan, but that the plan wasn't achieved in reality. This is most likely due to slippage in the schedule, as discussed in Chapter 5. The mine was operating with low inventories due to the struggle to achieve coal targets and, therefore, slippage in the schedule was much more likely to occur. In addition, two pits were approaching each other, causing a greater level of interaction between tasks than usual; for example, blasting in one pit would result in equipment in both pits having to park up during the blast.
2. The second observation relates to the difference between prime waste moved and coal mined. Prime waste excavation finished at 95% of plan, while coal mined finished at only 91% of plan. The lower percentage of coal mined should have resulted in an increase in coal inventory, but the coal inventory was unchanged, subsequently meaning that actual coal recovery was lower than that planned for in the Quarterly Plans. This was not unexpected, mine plans have coal loss and dilution assumptions built into them, but it is very common for mines to not have a good understanding of their real coal recovery and for them to use optimistic recovery assumptions in their planning. So this mine was only achieving 96% (the percentage of plan coal mined divided by the average percentage of planned for waste excavation, which is $91\% / 95\%$) of their planned recovery.

My advice to the mine was that, based on analysis of their historical performance against planned, they were not going to achieve the planned coal target for the following month and that the process to arrive at an estimate of what they would achieve was a fairly simple exercise.

Based on history, the dragline would achieve about 96% of the upcoming Quarterly Plan, so take the dragline equipment paths and drop the last 4% of the volumes moved during the period, any coal uncovered within that 4% was not going to happen and should be excluded from the estimate. The truck and shovel operations

also uncover coal, historically they would achieve about 94% of planned, so take the last 6% of waste moved in the period and if that uncovered any coal, it also could be removed from the schedule. Adding those two values together provided a new coal uncovered estimate, however, the mine was not recovering 100% of the planned coal. So, take the new coal uncovered value and multiply that by 96%, which is the real coal recovery historical performance and that will provide a coal mined value that I would have greater confidence in.

The example above highlights how it is possible to build reference class forecasting methodology into mine planning. I think there is always a need to check our complicated mine plans (built on a massive range of inputs) back against reality, something I like to refer to as “sanity checks”.

14

That's A Wrap!

Most reports and papers have an Executive Summary which is normally located at the beginning of the report. This is so you only have to read this summary to attain an understanding of what the paper is about, without having to read the whole paper. For this book, I created an Executive Summary (this chapter), but I placed it at the back, instead of the front. The reason I located it at the back of the book is that I don't want you to just read a summary. Because the devil is in the detail, so I want you to read the whole book and to follow the stories and examples used, which add weight to the arguments. I really do want to provoke the thinking that hopefully comes from some of the stories, examples and my comments.

I've created a table that summarises the greatest mine planning crimes from my experience in Table 14 - 1, to allow for easy reference. My recommendation is that you review the crimes that I've listed and put them in order of largest to smallest for your mine site. But the one thing that I implore is that you don't meet the definition of insanity - don't continue to do the same thing as you have always done and expect different results. You've gone to the effort to read this book. No mine site is perfect so I know that your mine site suffers at least one or more of the crimes that I have discussed in this book. So please, take action and get the ball rolling on fixing those problems.

My suggestion would be to review the issues that I have discussed in this book and whether they are relevant for your site. Add to those issues others you come up with for your site and then plot that complete list on an effort-return matrix as shown in Figure 10 - 1. That process will highlight for your mine site where you will get the greatest return with the least effort and which issues you might place at the bottom of the list.

Change is difficult; let's be honest, there are a lot of people that don't like change. That leads to a state called inertia, which is a tendency for the status to always return to the norm, that is the state before change was attempted. This is exemplified by the story I related in Chapter 1 about the one dragline mine where we did a lot of work and improved the dragline operation significantly, only for it to return to its previous performance once we removed the focus.

Chapter	The Crime	Recommendations
4 - Deterministic Scheduling	The continuing misguided belief in deterministic mine plans	Implement a stochastic mine planning solution (if you can find it), if you can't then....
		Apply pressure to software suppliers for stochastic capability
		Acknowledge the extremely high degree of variability within mine plans and change to working in date ranges rather than specific dates
		Manually calculate the probability of overlap for critical tasks
		Evaluate options for reducing the variability within tasks
13 - Optimistic Mine Planning	Using the mine plan to drive performance improvement	Separate mine planning process from the benchmarking process
	Lack of understanding of the impact of using optimistic assumptions	Calculate the probability of the schedule finishing ahead of plan for varying input assumption probabilities
	Producing mine plans that are unrealistic and have a short shelf life	Introduce a form of reference class forecasting that is applicable to your mine site as a check on schedule outputs
		Stop carrying out activities such as optimisation or destination scheduling, on optimistic plans - it's a waste of time
5 - Schedule Slippage	Lack of consideration of the slippage in schedules	For short term and execution plans, after calculating the probability of overlap for critical tasks, assess the likely degree of slippage in the mine plan
	Lack of understanding of the impact of skewed distributions	Carry out statistical analysis to develop an understanding of the degree of skew for each equipment or activity type
6 - Execution of the Plan	Unaware exactly what is important in a mine plan and therefore how to execute it	Select the critical interactions to focus on during execution of the plan
		Calculate the minimum/maximum shift targets required to ensure overlaps are avoided at critical interactions
		Change the weekly planning and communication to be based around management of the critical interactions
7 - Mine Planning KPI's	No KPI's that measure the quality of the mine plan and the mine planning process	Implement at least one KPI that measures the quality of the plan or planning process (two suggested, but there are others)
		Monitor the new KPIs and track trends to feedback into improving the planners and the process
		Stop measuring compliance to plan on a spatial basis - it is a waste of time
8 - Due Diligence	Incorrect plans due to mistakes within the input data or schedule itself	Allow the mine planner time to carry out the necessary checks before the plan is final
		Implement plan checking processes, e.g. checklists or peer reviews
9 - Understand Enablers	Not understanding how critical enablers are, so not being aware of, or scheduling them	Using the criteria provided, identify the enablers at your mine site
		Ensure that critical enablers are scheduled, or at least monitored to check if they are an issue
10 - Production Team Don't Follow The Plan	Those charged with executing the plan have no faith in it and prefer to implement their own plan	Get introspective, review your planning process for flaws
		Use an effort-return style matrix to determine which issues you will tackle first
		Ensure the right champion is engaged to build sustainable change
11 - Schedule The Right Activities	Scheduling too many activities or items of equipment, or worse still, not scheduling those that should be scheduled	Review long term mine planning process, can it be simplified? Freeing up labour resources for use elsewhere
		Review short term and execution plans, is all equipment being scheduled? In particular drill and blast and critical ancillary equipment
		If destination scheduling or haulage modelling, critically review whether it is required in each plan
12 - Understand Theory of Constraints	Not utilising any part of TOC within your mining processes	If you don't understand TOC, read up on it and build a basic understanding
		Identify the constraint at your mine
		Elevate the constraint and subordinate other activities to it within your mine planning

Table 14 - 1 Summary of the Crimes

Consequently, my recommendation is to choose a small effort item to tackle first, ideally something that is of high return. If you have some items that appear promising, don't make a big deal about it, don't make big announcements about a new process you're introducing. Employees have heard it all before and can be a bit cynical about change, so I would recommend quietly choosing the item you're going to tackle and Just Do It (borrowed from Nike). I suggest you get some success under the belt and some runs on the board before there is any fanfare.

Change needs a champion, so success is definitely dependent on engaging the right people. I think it is well worth reading the book "Influencer: The Power To Change Anything" that I referred to in Chapter 1. I recommend doing that before selecting people for the team, as the book has some good examples regarding the selection of certain types of personalities to drive change.

I've seen a lot of change successfully implemented in my time in the industry, but I haven't seen a lot of change successfully implemented on a sustainable basis. So don't just think about who is going to champion the change, but also who is going to champion sustaining the new practices, which may not be the person used as the change agent in the first place. I think someone to sustain the change is just as important as the person who makes the initial change.

There will be some very high effort items that will lead to very high returns that will be very tempting to implement first. But please don't start with those items, get sustained change on smaller items under the belt first and embed those changes, allowing you to learn along the journey, before tackling the bigger issues that require significantly more labour, co-ordination and effort to achieve success. Borrow from the Chinese proverb, "A journey of a thousand miles begins with a single step," and just take that first step, regardless of how small it is. Draw a leaf out of the Nike book and "Just Do It".

After writing this book, I've reflected on what was my number one takeaway. It was the ridiculous notion that we can carry out a deterministic mine plan, have an unwavering belief in it being right and then hold people accountable to follow it to the letter. I go back to the Peter Drucker quote that I opened this book with.

Trying to predict the future is like trying to drive down a country road at night with no lights while looking out the back window.

This quote is apt because we do look out the back window while trying to predict the future. We put a lot of effort into measuring data via fleet management systems so that we can build solid historical results for productivities and time usage. That is our “looking out the back window” that we then place great faith in being sufficient to construct a quality mine plan (predict the future).

I think the problem stems from the fact we call it a plan, which implies that it is something we can think about, plan ahead and then make happen – like it is easy. But it is not something we can just plan ahead and make happen, as I pointed out in Chapter 4, there is an extremely high degree of variability and, therefore, uncertainty in a mine plan. I believe part of the problem lies in our choice of language. I suspect we are all familiar with the term “variability” and comfortable with using this term, but if we used the term “uncertainty” instead of variability when discussing mine planning, that might help in dismissing our unfounded belief that we produce accurate and reliable mine plans. (I would note that in the world of statistics, the term uncertainty is a direct substitute for variability, they are interchangeable.)

For example, if instead of discussing the extremely high degree of variability in mine planning, I instead referred to the extremely high degree of uncertainty in mine plans, the phrase “extremely high degree of uncertainty” would lead me to having low confidence that it was actually going to happen as planned. Following directly on from this, I suggest we look at different terminology for “mine plan”, which by its very nature implies some reliability and confidence in the plan. What if instead of calling it a mine plan, we called it a mine prediction, or a mine estimate or mine projection? I know I would then think of it more as just a good guess as to what would happen at the mine site in the future.

What was your key takeaway from this book? This would be invaluable information to allow me to tailor my future articles and books to customers’ needs, so I would be very excited to hear any feedback you want to provide. My email address is info@markbowater.com for any comments.

I’ve thought about what I would define as success following the exceptional amount of time I have invested in writing this book. It would be awesome if it ignited change, such as the creation of stochastic scheduling tools, or the reduction in optimistic assumptions used in mine planning. But I need to be realistic, they are

substantial steps to take and I don't expect them to happen just because I said they should in this book.

I will feel the book has been a success if I can ignite discussion of mine planning issues outside of my immediate circle of contacts. Breaking the definition of insanity and eliminating these crimes requires action from those who are involved with mine planning.

The first step in fixing any problem is actually acknowledging that there is a problem in the first place. That's where the discussion needs to start across the industry, acknowledging that we have problems in mine planning and developing potential solutions. As I said in the Introduction this book contains my thoughts on ten mine planning crimes, but I think that is a good place for discussion to start, what do you see as the top mine planning crimes?

About 12 months back I created a LinkedIn group for discussion on mining engineering issues (with no advertising allowed), you can find that if you search for Miners Digs. It has only been of limited success in terms of facilitating the type of discussion I'd like to see. So if you have a suggestion as to how I can facilitate further discussion, I'm all ears and please let me know.

I see this book as an awesome opportunity for a simple survey of mine planning issues on a global scale so it would be awesome to hear your thoughts. I've created a simple and anonymous survey, which also has an area for comments. This survey provides an opportunity to understand the real issues within mine planning and the results will be completely transparent. Not only will every person who completes the survey receive a summary of the results, but they can provide their own questions for me to add for future survey respondents. I hope the insights gained by everyone from this survey will lead to genuine change in the industry. Link to the survey <https://www.surveymonkey.com/r/crimesagainstmineplanning>

I am also hoping the survey results and any other feedback from readers will provide subjects for future blog articles that I can write. I don't want to waste my time writing articles that are of little interest or no value to the mining industry. I want my articles to have an impact on the industry and the best way for me to do that is to know what the areas and subjects of interest and concern are. So please let me know via the email address previously provided, commenting on my blog at

<https://markbowater.com/> or contacting me through my LinkedIn profile <https://www.linkedin.com/in/mineplanningguru>

I will continue to post articles on LinkedIn and on my web site. These articles will cover a range of areas, including further thoughts on issues discussed in this book and further thoughts on areas of mining and mining engineering in general.

I would like to finish this book by reiterating that I am very keen to see change in the mining industry. My goal is to leave mine planning in a better place than I found it, which means that I'm keen to help facilitate that change. If you need help, please let me know. This is not about selling my services, but more about gaining an understanding of the needs of the industry and what is required for mine planning to move to a better place. For example, do we need new software tools, do we need better training or is better dissemination of information required throughout the industry? I can also write articles targeted to your needs or issues, that then will likely be of benefit to others.

Finally, please let me know of your successes and failures, they will provide valuable learning for both myself and others. Together, we can facilitate change that has been required for a long time.....

Thank You!

Epilogue - The Broken Stocks Saga

I think that stories provide invaluable opportunities to learn from and they are great for getting the brain ticking in terms of how that story might be applicable to you. One objective I have for this book is to get people thinking about changes they might make to mine planning and so I wanted to include plenty of stories about my mine planning experiences. The story that follows is one of my favourites and was pivotal in my journey down the path of challenging paradigms in the mine planning world and always looking for a better way.

I couldn't find a good home for it in any of the particular chapters as it relates to inventory management and (although I easily could have) I didn't include inventory management as one of my ten greatest crimes. But I wanted to include this story in the book as I think it is a great opportunity to learn from and to facilitate thinking as to whether to any elements from it are applicable to you. So I've added it as a close-out story to the book, one last piece to get you thinking.....

The Broken Stocks Saga

For a period in the '90s, I was Grade Control Superintendent at an iron ore mine in West Australia. The team I led was charged with producing weekly stockpiles of lump and fines ore. Between these two stockpiles, we had a total of eight quality targets with very tight bands that we had to achieve. Broken stocks (material already blasted) are critical in this weekly ore stockpile construction as they're typically all there is available to mine from within that time frame.

The Mine Planning Department was responsible for short term scheduling of all mining equipment over a three-month time frame. In that plan, while drill, blast and excavation were all scheduled, the primary driver was sequencing the excavators to mine the quantity and quality of product required.

Not long after moving in to the Grade Control Superintendent role, we came into a situation where we had very high phosphorus in our broken stocks, which made it nearly impossible to build weekly stockpiles that met our quality targets. The problem was that the Mine Planning Department, in scheduling the diggers to mine product to specification, were effectively scheduling specification product to be

added at the end of the schedule. However, every week, our role was to take a specification batch away from the inventory. So, the inventory (broken stocks) was not being corrected and for a sustained period of nearly two months, the broken stocks continued at high levels of phosphorus. During that entire period, (about nine stockpile builds) we really struggled to produce specification weekly stockpiles. This was very much a saga for me, as I drew a lot of heat for railing off-spec product.

This situation was a catalyst for many learnings, the first of which was about broken stocks management. My first learning was how critical broken stocks are to achieving target qualities for each stockpile construction, which cannot be understated. If the Grade Controllers are provided with a set of broken stocks that are close to quality specs, it is relatively easy for them to hit their target. On the other hand, if they are provided with broken stocks which have one or more qualities a long way from target, it is extremely difficult.

Similarly, provide a large quantity of broken stocks and you provide more choice, provide a small quantity and, even though they might average out at the required specification, the lack of choice (and therefore scheduling flexibility) may still make it difficult to achieve target.

While the short term planning group were scheduling the shovel excavation for the next month, I believe a much better solution would have been for them to manage the broken stocks instead. The short term mine planning process should have involved determining what the starting broken stocks were (quantity and quality) and then assessing which blocks need to be added to this to either maintain stocks within target bands or restore them to within target bands.

In this way, scheduling is effectively an inventory management process which is both a critical part of mine scheduling and one I see as typically underutilised.

Recommended Reading (& Bibliography)

Some of the books below I have referred to in the text, some I haven't. But I would recommend all of them as reading that have nothing to do with mine planning directly, but indirectly they are of great benefit.

Flyvbjerg, B. (2007). Curbing Optimism Bias and Strategic Misrepresentation in Planning: Reference Class Forecasting in Practice. *European Planning Studies*, 16(1), 3–21. <https://doi.org/10.1080/09654310701747936>

Gawande, A. (2009). *The Checklist Manifesto: How to Get Things Right* (1st ed.). Metropolitan Books.

Goldratt, E. M. (2002). *Critical Chain*. North River Press.

Goldratt, E. M. (2014). *The Goal* (Fourth Revised (30th Anniversary) Edition). North River Press.

Grenny, J., Patterson, K., Maxfield, D., McMillan, R., & Switzler, A. (2013). *Influencer: The New Science of Leading Change, Second Edition* (2nd ed.). McGraw Hill.

Kerr, J., & Reichlin, S. (2016). *Legacy (Sports)* (Unabridged ed.). Audible Studios on Brilliance Audio.

Koontz, D. (2006). *VELOCITY*. Bantam Books.

Savage, S. L., Danziger, J., & Markowitz, H. M. (2012). *The Flaw of Averages: Why We Underestimate Risk in the Face of Uncertainty* (1st ed.). Wiley.

Taleb, N. N. (2005). *Fooled by Randomness: The Hidden Role of Chance in Life and in the Markets (Incerto)* (2nd ed.). Random House Trade Paperbacks.

Taleb, N. N. (2007). *The Black Swan: The Impact of the Highly Improbable (Incerto)* (1st ed.). Random House.

Martyn Bloss, Geoff Capes, Russell Seib, Liam Alford, Jack Light, Ilnur Minniakhmetov and Chris Nielsen (2020). *Value Chain Excellence - managing variability to stabilise and exploit the mine value chain*.

'Mark courageously paces you through a career of learning with insightful, honest and enjoyable scenarios and stories. I would recommend you get reading and have your notepad ready'

Andrew McDougall
Technical Head – Mining

CHANGING MINE PLANNING FOREVER

In this ground-breaking book, Mark, with more than 30 years experience in mine planning, using stories and examples, clearly lays out what he believes to be the ten worst crimes against mine planning still commonplace today. With insight on not just day to day planning but the system as a whole, Mark provides practical advice on correcting, avoiding, and overcoming each issue.

The book covers the problems with deterministic mine planning, schedule slippage, mine planning KPIs, optimistic plan inputs, how to implement the plan, and frustration from other teams about the reliability of the mine plan, to name a few. This book will provide you with the practical tools to not just recognise the issues at your mine site, but also to improve your schedules, your mine processes and your team dynamics and relationships.

This book isn't just for mine planners, it is for all stakeholders of the mine planning process and particularly for those in positions of influence who have the authority to help change planning at their mine site for the better. This book is for anyone who has ever turned up to work in the mining industry and thought, there's got to be a better way to do this.



Mark Bowater has been involved in the Australian open cut mining industry for over thirty years, that includes working for a major mining company, large scale contractor and running his own consulting business for over 15 years. While he has carried out a wide range of roles and studies, he has always been passionate about mine planning and the dichotomy of it being so heavily relied upon in the industry, while also being so widely scorned for its inadequacy. Mark has set himself a legacy to leave mine planning in a better place than he found it. This book is part of that purpose.

ISBN: 978-0-6454098-2-6