UNIVERSITY OF TORONTO | Engineering

Reinforcement learning approach to finite inventory problem

1910s to 1930s: The Dawn of Inventory Theory

1913 first publication of EOQ model
by a Production Engineer Ford Harris

- Practical approach and simple use
- Quickly adopted by the industry
- Erroneously cited up until 1988

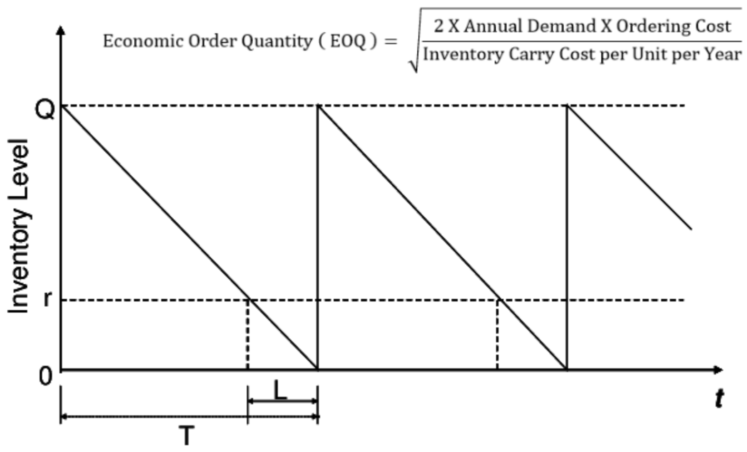

Image source: Herrera, Ozdemir, Rıos (2009) Capacity planning in a telecommunications network: a case study. Int J Ind Eng

HOW MANY PARTS TO MAKE AT ONCE

FORD W. HARRIS
Production Engineer

Reprinted from *Factory, The Magazine of Management*, Volume 10, Number 2, February 1913, pp. 135–136, 152

Interest on capital tied up in wages, material and overhead sets a maximum limit to the quantity of parts which can be profitably manufactured at one time; "set-up" costs on the job fix the minimum. Experience has shown one manager a way to determine the economical size of lots.

Every manufacturer is confronted with the problem of finding the most economical quantity to manufacture in putting through an order. This is a general problem and admits of a general solution, and, however much it may be advisable to exercise judgment in a particular case, such exercise of judgment will be assisted by a knowledge of the general solution.

The writer has seen the practical workings of a first-class stock system and does not wish to be understood as claiming that any mere mathematical formula should be depended upon entirely for determining the amount of stock that should be carried or put through on an order. This is a matter that calls, in each case, for a trained judgment, for which there is no substitute. There are many other factors of even more importance than those given in this discussion.

But in deciding on the best size of order, the man responsible should consider all the factors that are mentioned. While it is perfectly possible to estimate closely enough what effect these factors will have, the chances are many mistakes costing money will be made. Hence, using the formula as a check, is at least warranted. Given the theoretically correct result, it is easy to apply such correction factors as may be deemed necessary.

In determining the economical size of lot the following factors are involved:

Unit Cost (*C*). This is the cost in dollars per unit of output under continuous production, without considering the set-up or getting-ready expense, or the cost of carrying the stock after it is made.

Set-up Cost (*S*). This involves more than the cost of getting the materials and tools ready to start work on an order. It involves also, the cost of handling the order in the office and throughout the factory. This cost is often neglected in considering the question. Most managers, indeed, have a rather hazy idea as to just what this cost amounts to. If such is the case an investigation will show that the cost of handling, checking, indexing and superintending an order in the offices and shops is a considerable item and may, in a large factory, exceed *one dollar* per order.

The set-up cost proper is generally understood. Indeed, shop foremen in general appreciate only too well what the cost of set-up means on small orders, and so, if left to themselves, will almost invariably put their work through in large quantities to keep down this item. So doing, however, affects unfavorably the next factor.

Interest and Depreciation on Stock (*I*). Large orders in the shop mean large deliveries to the storeroom, and large deliveries mean carrying a large stock. Carrying a large stock means a lot of money tied up and a heavy depreciation. It will here be assumed that a charge of ten per cent on stock is a fair one to cover both interest and depreciation. It is probable that double this would be fairer in many instances.

Movement (*M*). It is evident that the greater the movement of the stock the larger can be the quantities manufactured on an order. This, then, is a *vital* factor.

Manufacturing Interval (*T*). This is the time required to make up and deliver to the storeroom an order, and, while it seldom is a vital factor, it is of value in the discussion.

There is another factor, *X*, the *unknown* size of order which will be most economical. Thus summarizing, there are the following factors in the problem:

M equals the number of units used per month (movement).
C equals the quantity cost of a unit in dollars or the *unit* cost.

1950s to 1960s: The Golden Era of Inventory Theory

US military funded research and "Stanford Studies"

- Foundations of Stochastic Lot Sizing Problems (Q,R); (s,S) policies
- Wagner Whitin *dynamic programming* algorithm to solve Classical Dynamic Lot Sizing Problem
- First major textbooks on inventory control

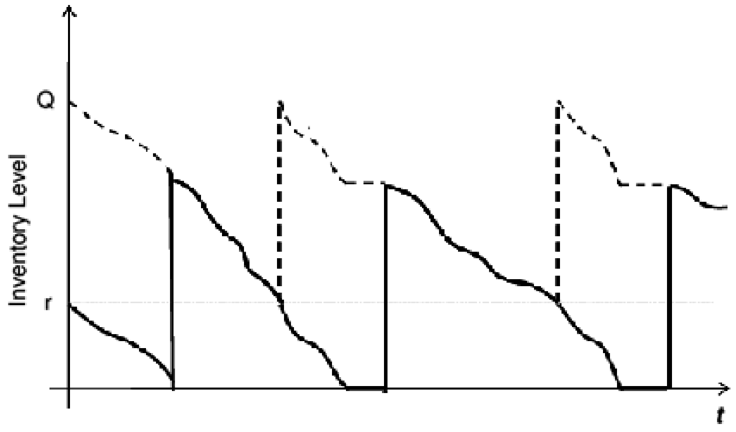

Image source: Comez, Kiessling, (2012). Joint inventory and constant price decisions for a continuous review system

Image source: https://www.flickr.com/photos/123727295@N07/14498214658, by T. Schroeder, cropped. Free to share with attribution

1950s to 1960s: The Golden Era of Inventory Theory

US military funded research and “Stanford Studies”

- Foundations of Stochastic Lot Sizing Problems (Q,R); (s,S) policies
- Wagner Whitin *dynamic programming* algorithm to solve Classical Dynamic Lot Sizing Problem
- First major textbooks on inventory control

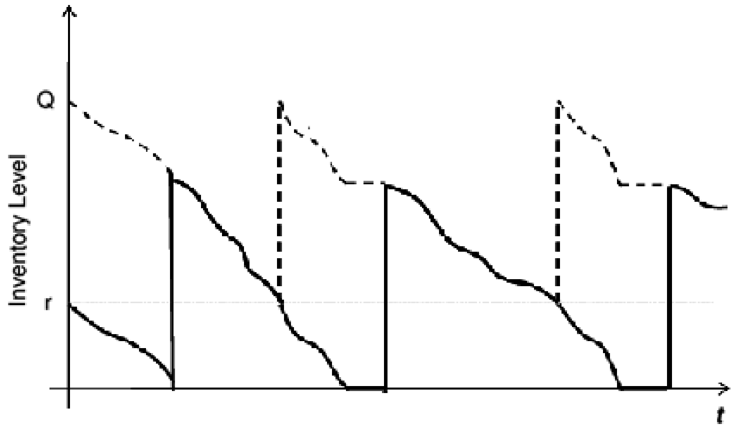

Image source: Comez, Kiessling, (2012). Joint inventory and constant price decisions for a continuous review system

1970s:

- Industry often prefers classic EOQ Challenges with adoption
- Limited computational power
- “The natural academic drift” in Operations Research

1980s to 1990s: The Digital Era of Inventory Theory

Personal computers, local networks and the Internet

- Heuristic and simulation models
- Industry adoption
- New challenges – the data

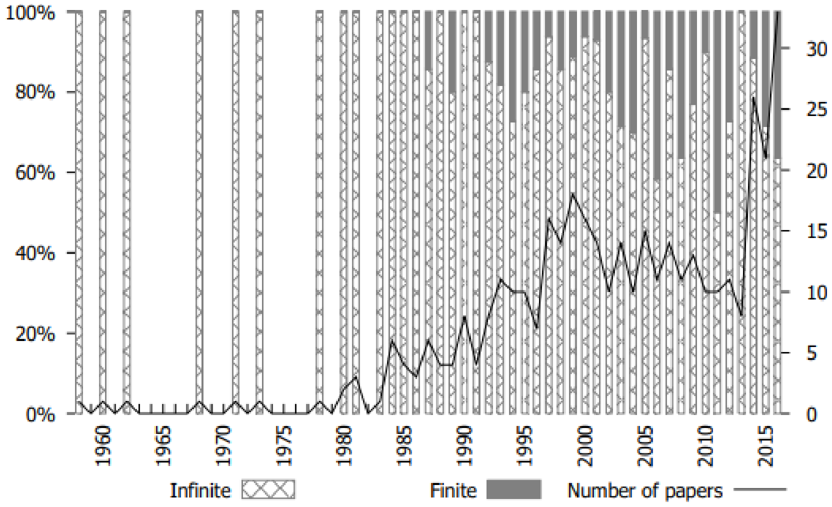

Image source: de Kok, Grob, Laumanns, Minner, Rambau, Schade, (2018). A typology and literature review on stochastic multi-echelon inventory models (Figure 6: Capacity)

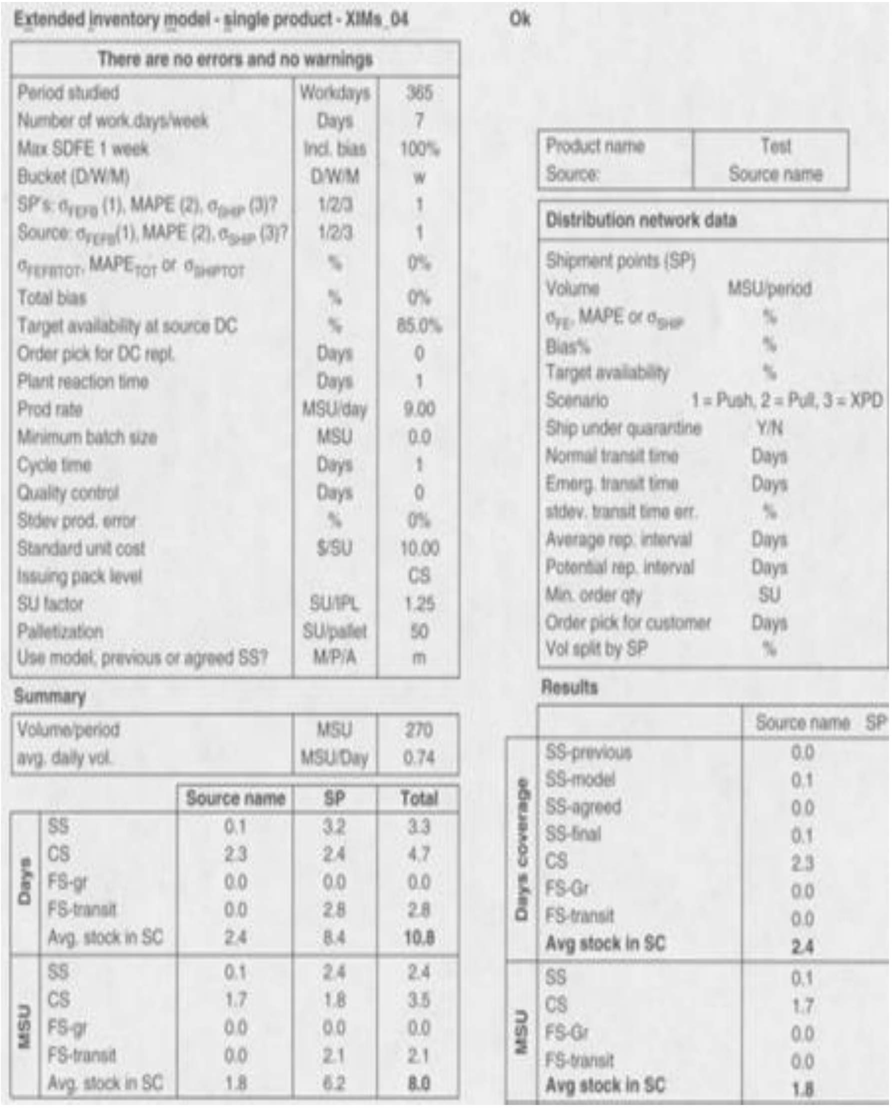

Extended inventory model - single product - XIMs_04

Ok

There are no errors and no warnings		
Period studied	Workdays	365
Number of work.days/week	Days	7
Max SDFE 1 week	Incl. bias	100%
Bucket (D/W/M)	D/W/M	w
SP's: σ_{FEFB} (1), MAPE (2), σ_{SHIP} (3)?	1/2/3	1
Source: σ_{FEFB}(1), MAPE (2), σ_{SHIP} (3)?	1/2/3	1
$\sigma_{FEFBTOT}$, $MAPE_{TOT}$ or $\sigma_{SHIPTOT}$	%	0%
Total bias	%	0%
Target availability at source DC	%	85.0%
Order pick for DC repl.	Days	0
Plant reaction time	Days	1
Prod rate	MSU/day	9.00
Minimum batch size	MSU	0.0
Cycle time	Days	1
Quality control	Days	0
Stdev prod. error	%	0%
Standard unit cost	$/SU	10.00
Issuing pack level		CS
SU factor	SU/IPL	1.25
Palletization	SU/pallet	50
Use model, previous or agreed SS?	M/P/A	m

Summary

Volume/period	MSU	270
avg. daily vol.	MSU/Day	0.74

		Source name	SP	Total
Days	SS	0.1	3.2	3.3
	CS	2.3	2.4	4.7
	FS-gr	0.0	0.0	0.0
	FS-transit	0.0	2.8	2.8
	Avg. stock in SC	2.4	8.4	**10.8**
MSU	SS	0.1	2.4	2.4
	CS	1.7	1.8	3.5
	FS-gr	0.0	0.0	0.0
	FS-transit	0.0	2.1	2.1
	Avg. stock in SC	1.8	6.2	**8.0**

Product name	Test
Source:	Source name

Distribution network data	
Shipment points (SP)	
Volume	MSU/period
σ_{FE}, MAPE or σ_{SHIP}	%
Bias%	%
Target availability	%
Scenario	1 = Push, 2 = Pull, 3 = XPD
Ship under quarantine	Y/N
Normal transit time	Days
Emerg. transit time	Days
stdev. transit time err.	%
Average rep. interval	Days
Potential rep. interval	Days
Min. order qty	SU
Order pick for customer	Days
Vol split by SP	%

Results

		Source name	SP
Days coverage	SS-previous	0.0	
	SS-model	0.1	
	SS-agreed	0.0	
	SS-final	0.1	
	CS	2.3	
	FS-Gr	0.0	
	FS-transit	0.0	
	Avg stock in SC	**2.4**	
MSU	SS	0.1	
	CS	1.7	
	FS-Gr	0.0	
	FS-transit	0.0	
	Avg stock in SC	**1.8**	

Image source: Farasyn, Perkoz, Van de Velde, (2008). Spreadsheet Models for Inventory Target Setting at Procter & Gamble (Figure 1: The screenshot shows a sample XIM model screen). Cropped

Starting 2010s: The AI Era of Inventory Theory?

Neural Networks

- Warren McCulloch and Walter Pitts proposed computational model of the “nerve net” as early as 1943 [1]
- In 1986 Geoffrey Hinton (UofT), David Rumelhart, and Ronald Williams published back propagation algorithm for training [2]
- The decade 2010s… [3]

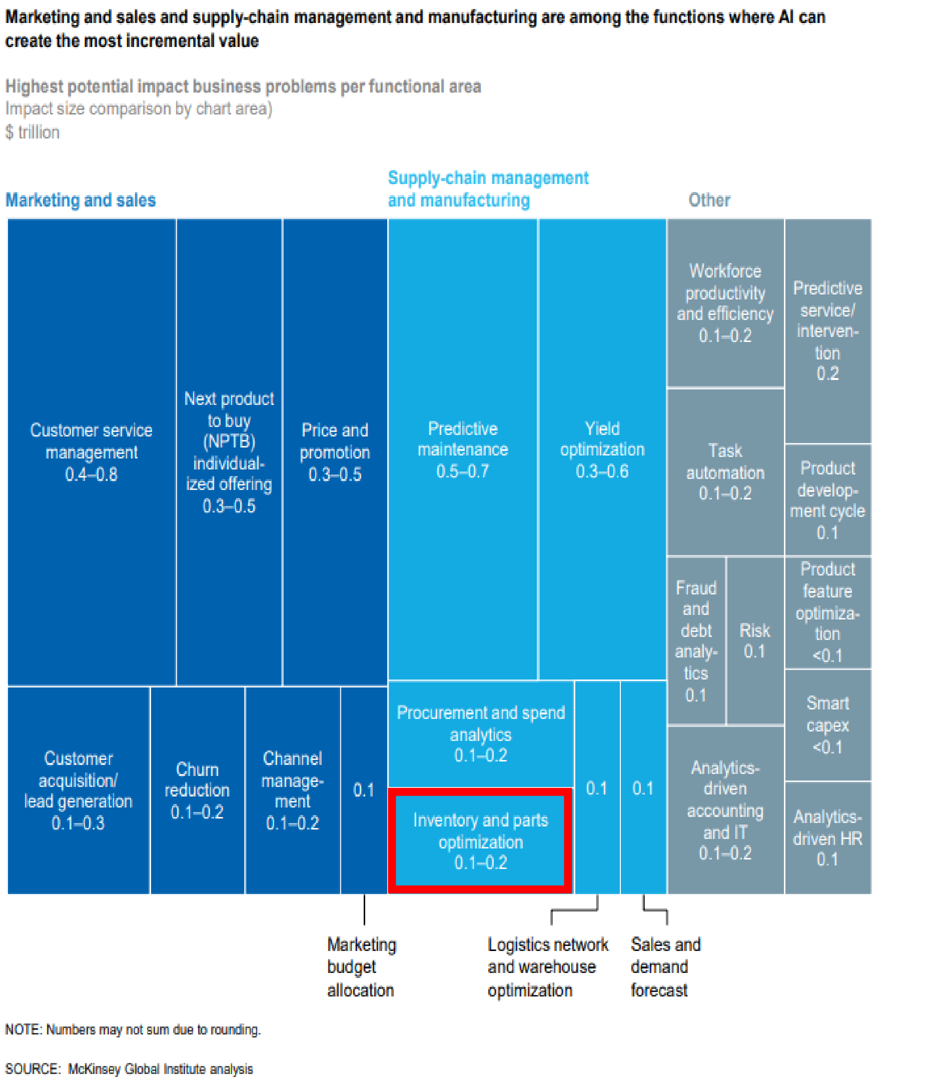

Image source: Chui, Manyika, Miremadi, Henke, Chung, Nel, Malhotra.
Notes from the AI frontier: Insights from hundreds of use cases. Discussion paper. April 2018. Exhibit 12, p. 21

References:

[1] W. McCulloch and W. Pitts, 1943. A logical calculus of the ideas immanent in nervous activity. Bulletin of Mathematical Biophysics, vol. 5
[2] D. Rumelhart, G. Hinton, R. Wiliams, 1986. Learning representations by back-propagating errors. Nature, vol. 323
[3] The Economist, Special Report on AI in Business, March 2018

Reinforcement Learning

Where does the field of Reinforcement Learning fit?

- Reinforcement Learning is solving MDP
- Dynamic Programming, Monte Carlo, Temporal Difference methods
- Approximate Solution Methods

Image source: https://en.wikipedia.org/wiki/Breakout_(video_game)

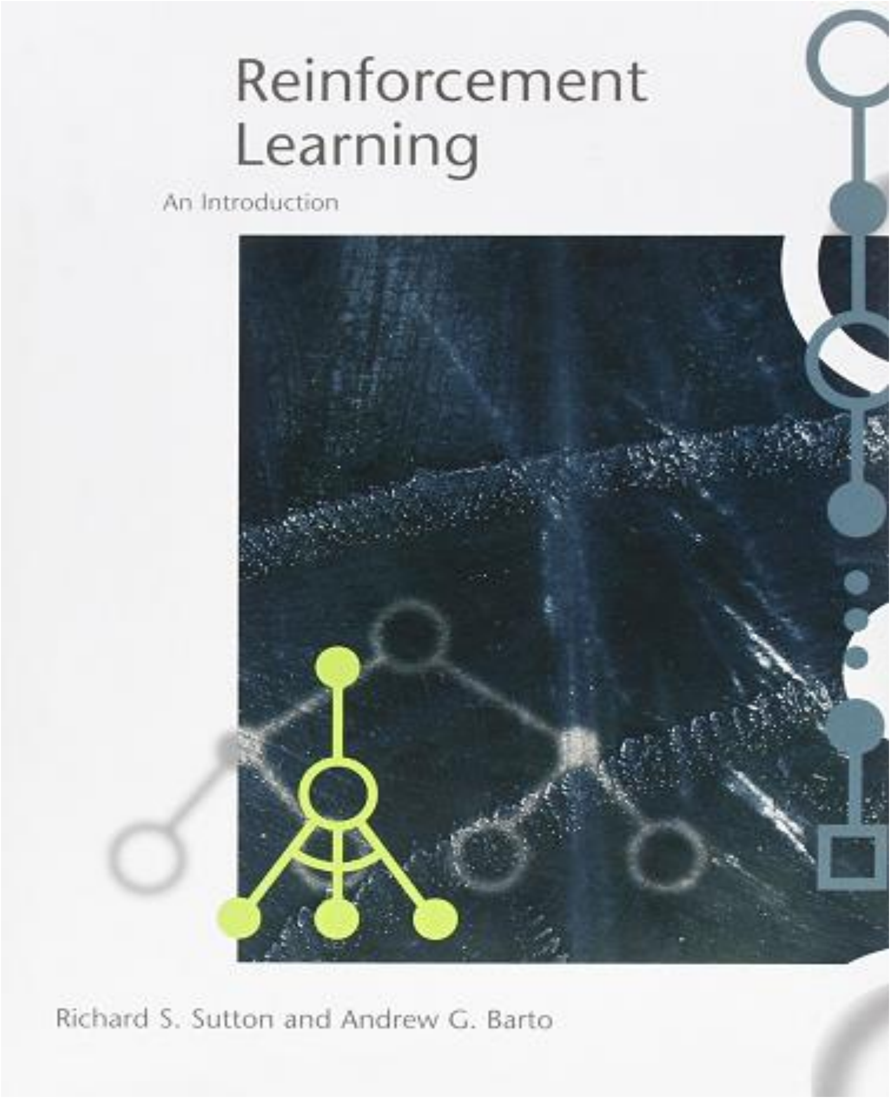

Image source: Sutton, Barto. Introduction to Reinforcement Learning, MIT Press Cambridge, MA, USA – book cover

Proof of concept models

Deep Q Learning algorithm
(Off-Policy Temporal Difference Control)

- 1 item: 0.04% less than optimal solution
- 2 items: buying sub-optimally each item saved 18% on simultaneous procurement

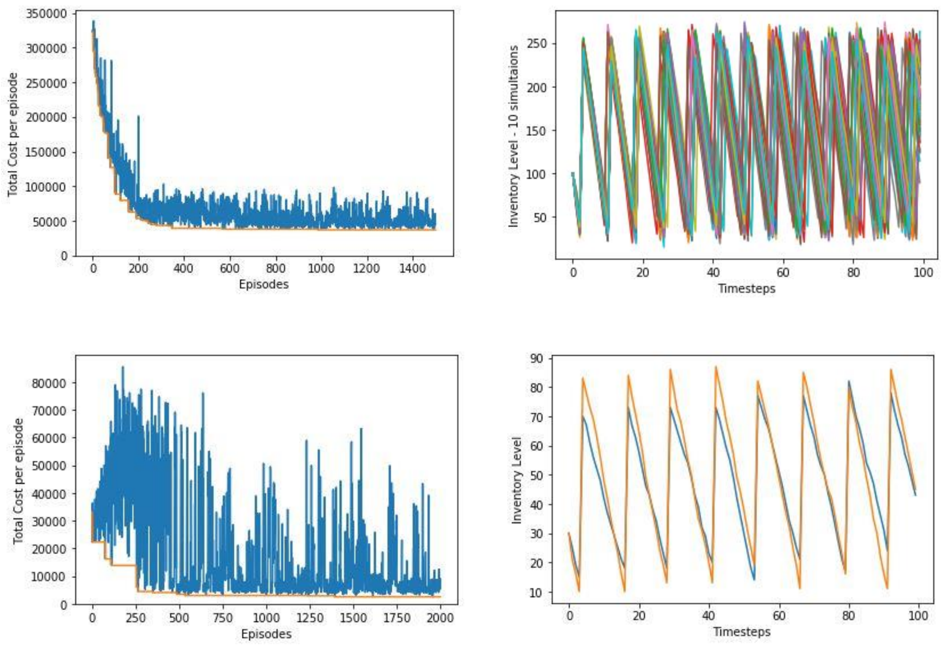

Policy Gradient algorithm

- 1 item: Converged to near-optimal solution <1% less than optimal solution

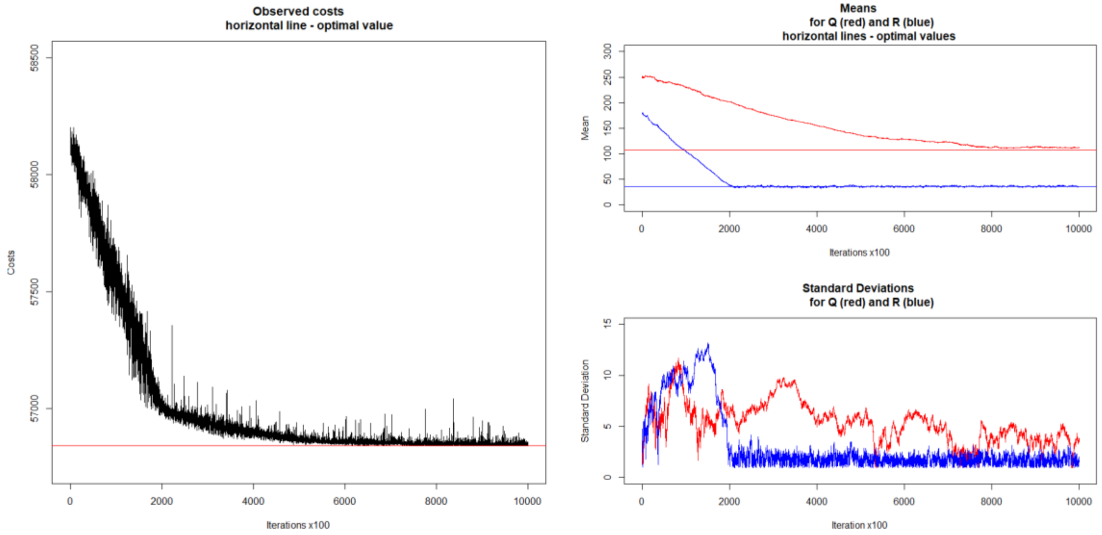

Thank you!