

THE LIMITATIONS OF SUB-OPTIMAL POLICIES

Izak Duenyas
Department of Industrial and Operations Engineering
The University of Michigan
Ann Arbor, MI 48109-2117

Technical Report 93-8

February 1993

The Limitations of Sub-Optimal Policies

IZAK DUENYAS

Department of Industrial and Operations Engineering

The University of Michigan

Ann Arbor, Michigan 48109

Abstract

In a recent paper, Zangwill(1992) constructed an example where setup reduction increased inventory and cost. Zangwill used this example to argue that if set-up times are cut, inventory and costs can increase, and that this is in conflict with the principles of “Japanese production theory”. We argue that the increase in costs that Zangwill(1992) demonstrated is due to his use of the “exhaustive” polling policy, which is a suboptimal policy. Using Zangwill’s example, we show how a different policy for the same problem, recently suggested in Van Oyen and Duenyas (1992), significantly outperforms the “exhaustive” policy used in Zangwill (1992), and also results in decreasing inventory and costs when setup times are cut.

In a recent paper, Zangwill (1992) constructed an example where reduction of setup times led to an increase in inventory costs. In fact, in Zangwill’s example, reduction of set-up times results in inventory, and therefore the costs of carrying inventory, to grow without bound. Zangwill (1992) argued that these examples expose a flaw in the current “Japanese production theory”.

Zangwill’s example involves a machine that makes n different items. The machine makes an item from kits of supply parts that arrive following a Poisson distribution. The time to make item i from a kit is s_i , which is an independent and identically distributed random variable. If the machine is set up to produce items of type j , and management decides to switch to

producing items of type i , a set-up is required. The duration of this set-up is an independent, identically distributed random variable d_i , with mean D_i .

In Zangwill's example, the machine produces items in the following manner: First, the machine is set up to make item 1, then the machine keeps making item 1 until all kits for item 1 (including those that arrived after the machine was set up) are exhausted. Then, the machine is set-up for production of item 2 (regardless of whether a kit for item 2 is available at that point). Then, all kits for item 2 are exhausted, and the machine is set up for item 3 and so on. After finishing all kits for item n , the cycle starts again. Thus, the machine uses an "exhaustive polling policy" to process the different items.

We first note that if the purpose of the facility is to provide quick response times to customers, and to minimize inventory carrying costs for kits, the "exhaustive polling policy" described above is far from being an "optimal" policy. It is only a heuristic policy since there is no reason why the machine *should* be controlled in this particular manner. Furthermore, it is not even a reasonably good heuristic policy. To see this, assume that somehow the facility is able to reduce all set-up times to 0. In this case, it has been shown that the policy that minimizes inventory costs is the $c\mu$ rule, where c is the inventory cost rate for an item and $\mu = 1/Es$. (Buyukkoc et al., 1985). If the inventory costs are the same for all the products, then this rule results in the production of the item with the shortest mean processing time. Hence, if $Es_i < Es_j$, the machine should not keep producing items of type j when a kit for item i is available. However, it is clear that the "exhaustive polling policy" will keep on producing items of type j , until the kits for item j are exhausted. Hence, this policy can result in large inventory costs, *even if the facility is able to reduce set-up times to 0*.

The fact that using the "exhaustive polling policy" in this problem is a sub-optimal heuristic policy provides an explanation for the "paradoxes" described in Zangwill (1992). Sub-optimal heuristic solutions can display unexpected behavior. This behavior does not necessarily indicate a flaw in the underlying theory. In fact, it often indicates that there is something wrong with the sub-optimal heuristic policy that is being used. Consider a practitioner who has reduced

set-up times and found that inventory has grown. Should she think that there is something wrong with reducing set-up times and that maybe she should be *increasing* set-up times or should she perhaps be paying more attention to the particular heuristic production policy giving rise to this behavior?

Clearly, if the optimal production policy (i.e., that policy which results in production with the minimum possible cost, given the set-up times) were known, then one could implement the optimal policy and be able to observe whether the same behavior, inconsistent with the theory, can be observed. Unfortunately, for the problem described above, the optimal policy is not known in the case where set-up times are non-zero. However, recently researchers have partially characterized the characteristics of the optimal policy. (Hofri and Ross (1987), Liu et al (1992), Van Oyen and Duenyas (1992)). For example, Van Oyen and Duenyas (1992) showed that when the machine is set-up to produce the item with the highest index (where item i is given the index c_i/μ_i), it is indeed optimal to keep producing until all kits for that item are exhausted. However, it is clear that if the machine is set-up to produce items with the lower index, it may sometimes be optimal to switch to the production of an item with a higher index before the kits for the item being produced are exhausted. For example, in the case where the inventory carrying costs for all kits are identical, it may make sense to switch to producing the product with the shortest processing time, if the set-up times are “low enough”, and “enough” kits for the product with the higher index have accumulated. This is due to the fact that even though a set-up will be required at first, a high number of products can be produced quickly after the set-up. Hence, the large decrease in the total amount of inventory after the set-up is completed and production begins, might more than compensate for the initial time lost to the set-up.

Even though the optimal policy for a problem may not have been completely characterized, it may still be possible to derive an effective heuristic solution by making use of some of the properties that the optimal solution has been shown to have. In fact, for the problem described above, Van Oyen and Duenyas (1992) recently derived an effective heuristic policy

in precisely this manner by making use of known results about the properties of the optimal policy. In a comprehensive simulation study (reported in Van Oyen and Duenyas (1992)), this heuristic outperformed every other heuristic considered in the literature including the “exhaustive polling policy” for every example considered. We describe this heuristic for a facility that produces two different items, and then show, using the example by Zangwill, that this heuristic results in significantly better performance than the “exhaustive polling policy” described above. (A heuristic was also derived in Van Oyen and Duenyas for $n > 2$ items, however, for simplicity, we focus on the two item case.)

The heuristic policy derived in Van Oyen and Duenyas (1992) is an index policy that is very simple to implement. Without loss of generality, we let the item with the higher $c\mu$ index be item 1. We let x_i denote the number of kits of type i that are available, at the time when a decision is being made on what to produce next. As we noted above, if the machine is setup for production of item 1, and $x_1 > 0$, it is optimal to produce another item of product 1. Therefore, suppose that the machine is set-up to produce item 2, and $x_2 > 0$, and $x_1 > 0$. We define φ as

$$\varphi = c_1\mu_1 \frac{\frac{x_1 + \lambda_1 D_1}{\mu_1 - \lambda_1}}{D_1 + \frac{x_1 + \lambda_1 D_1}{\mu_1 - \lambda_1} + D_2}. \quad (1)$$

where λ_i is the arrival rate of type i kits. In this case, the policy is to setup the machine and start production of item 1 if $\varphi > \rho c_1\mu_1 + (1 - \rho)c_2\mu_2$, where $\rho = \lambda_1/\mu_1 + \lambda_2/\mu_2$. Otherwise, one more unit of product 2 should be produced. To complete the description of the policy, we also need to specify what the operator of the machine should do when the kits for the item the machine is set-up to produce are exhausted. Suppose the machine is set up for item 1, and $x_1 = 0$, then the policy is to set the machine up for production of item 2 if $x_2 > \lambda_2 D_1$. Otherwise, the machine is kept idle until the next arrival of a kit to the system. Similarly, if the machine is set up for production of item 2, and $x_2 = 0$, then the machine is currently set-up for production of item 1 only if $x_1 > \lambda_1 D_2$. Finally, the policy does not allow for a set-up unless at least one item has been completed since the last set-up. For the interested reader, we provide more details on the derivation of this heuristic in Appendix 1. A comprehensive

study of the performance of this heuristic can be found in Van Oyen and Duenyas (1992).

Having described an alternative heuristic policy to the exhaustive polling policy, we now return to the class of examples constructed by Zangwill where the exhaustive polling policy results in increased costs when set-up times are decreased. We consider a machine that produces two items. We assume that the inventory holding cost for the kits are identical, so that inventory costs are proportional to the average number of kits in the system (or to the expected waiting time of a kit). Notice that our heuristic can also be used if holding costs are different and this assumption is for expository purposes only. The set-up time for item 1 is deterministic and 0.01 hours. The set-up time for item 2, d_2 is $10/u^2$ hours with probability $1 - 1/u^3$, and $10u$ hours with probability $1/u^3$. (It is easy to check that as u goes to infinity, $Var(d_2)/D_2$ diverges to infinity, which, as Zangwill demonstrated is a sufficient condition for inventory to grow without bound when the exhaustive polling policy is used.) The Poisson arrival rate of kits for item 1 (item 2) is 0.08 (0.025) items/hr. The processing times are also deterministic and 4 hours for item 1 and 10 hours for item 2, respectively.

Exact results for waiting times under the exhaustive polling policy used in Zangwill (1992) were derived in Takagi (1986), and Sarkar and Zangwill (1991). We can use those results to compute the expected waiting time for an item of type 1 or type 2. To compute the average waiting times under our heuristic, described above, we used simulation. For each value of u , 20 independent runs were made, each consisting of 100000 item completions. The average waiting times for items 1 and 2 were separately computed.

We display the results we obtained from our simulation study, as well as the results for the exhaustive polling policy in Table 1. Notice that as u is increased from 2 to 7, the average set-up times, D_2 as well as the variance of the set-up times $Var(d_2)$ decrease. However, as Zangwill predicted, under the exhaustive polling policy, the average waiting times for item 1, $W_1(ex)$, and item 2, $W_2(ex)$, increase (which therefore imply, by Little's law that the average inventory is also increasing). $W(ex)$ denotes the average waiting time for a job. (Notice that $W(ex)$ is not the average of $W_1(ex)$ and $W_2(ex)$ since kits for item 1 and item 2 arrive at

different rates). The mean waiting times for item 1 ($W_1(h)$), item 2 ($W_2(h)$), and all items, ($W(h)$), under our heuristic policy (along with the standard deviation for the 20 runs) are also displayed in Table 1. Notice that, in all cases, the average waiting time is significantly lower than that under the exhaustive polling policy. Furthermore, for higher values of u , the mean and the variance of set-up times for item 2 decrease, and this results in the average waiting time for an item to decrease as well. Hence, under our heuristic, average inventory, and inventory costs decrease even though $Var(d_i)/D_i$ is diverging.

We also considered the case where set-up times for both items are 0. As we have previously noted, in this case, the $c\mu$ rule is optimal. Since the average waiting time when set-up times are 0 is a lower bound on the performance of our heuristic, we simulated this case and obtained a value of 9.11 for the average waiting time of a job. Hence, as u is increased, and thereby average set-up times are decreased, the waiting times under our heuristic are getting very close to the theoretically minimum possible average waiting time that can be obtained when all set-up times are 0.

The last column of Table 1 displays an interesting result. “Japanese inventory theory,” would predict that, in this example, it is better for the machine to process item 1 when kits for item 1 are available, since item 1 can be processed much faster, and wasteful inventory can be decreased more rapidly. However, when set-up times for item 2 are high, it is not possible to switch over to production of item 1 very rapidly. In Table 1, s^* denotes the minimum number of kits of type 1 that must be available before the machine is allowed to leave unfinished work of type 2 (i.e., $x_2 > 0$), and switch to item 1 (The values for s^* can be easily calculated from equation (1)). When $u = 2$, set-up times are high, and at least 4 kits for item 1 are required before **switching** is allowed. As u is increased, and thereby, average set-up times are decreased, the number of kits of type 1 required to allow switching decrease. Finally, when $u = 4$, the set-up times are low enough that the machine is allowed to switch to production of item 1 as soon as only one kit for item 1 is available.

It is interesting to note that in describing the guiding principles of Japanese inventory

u	D_2	Var_d_2	$W_1(ex)$	$W_2(ex)$	$W(ex)$	$W_1(h)$	$W_2(h)$	$W(h)$	s^*
2	7.031	75.4	16.26	14.67	15.88	10.58 (0.09)	21.64 (0.24)	13.21	4
3	3.272	67.0	18.88	15.30	18.02	8.38 (0.06)	20.41 (0.35)	11.24	2
4	1.860	53.7	22.53	17.41	21.31	7.50 (0.07)	19.74 (0.38)	10.41	1
7	0.611	31.8	34.56	25.40	32.38	7.12 (0.04)	17.76 (0.29)	9.65	1

Table 1: Comparison of Two Policies

theory, Zangwill writes (p. 15)

(...) Inventory is a reflection of waste, this theory proclaims, and the more the inventory, the more the underlying waste.(...)

(...) Reducing setups allows small production batches and low inventory, which cuts cost. Much more important perhaps, it facilitates mixed-model production. Fast setups enable a firm to make a small quantity of one model and then quickly switch over to make another model.

In our example, processing jobs of type 1, when kits for item 1 are available, is preferable because item 1 can be processed more rapidly and hence in a given interval, more “wasteful” inventory can be transformed into “product” than when processing item 2. It is not possible to switch to production of item 1 very quickly when set-up times are high. However, when set-up times are reduced, the firm gets the opportunity to switch to production of item 1 and to lower “wasteful” inventory more rapidly. As the s^* values indicate, our heuristic makes use of **this** increased capability to switch more rapidly, and decreases inventory and inventory costs, when set-up times are reduced. However, processing jobs according to an “exhaustive polling” policy even when set-up times are lower does not use this capability of mixed-model production to reduce wasteful inventory, and results in large amounts of inventory. Hence, Japanese production theory would indeed predict that the “exhaustive polling policy” is not a very good heuristic policy for this problem.

Conclusions

We have argued that paradoxes and logical inconsistencies observed when using a suboptimal heuristic solution should not lead to the immediate conclusion that the underlying theory is flawed. Rather, the particular suboptimal heuristic in use should be questioned first, since that process may lead one to finding the “optimal” solution or at least to be able to characterize some properties of the optimal solution in order to obtain a more effective heuristic. For the particular examples of paradoxes in Zangwill (1992), we have demonstrated that Japanese production theory would indicate that the particular suboptimal heuristic policy (i.e., the exhaustive polling policy) being used in these examples is not a good policy. We have also provided another heuristic policy for the same problem which greatly outperforms the exhaustive polling policy and has the behavior that Japanese production theory predicts.

Appendix 1

Our heuristic is based on the notion of “reward rates” developed in Van Oyen and Duenyas (1992). The machine is said to be earning rewards at a rate of $c_i\mu_i$ when processing jobs of type i . The quantity $c_i\mu_i$ can be regarded as the reward rate at which inventory carrying costs are reduced by processing a job of type i . In calculating an index for switching to type 1 production from type 2 production, we assume that when kits of type 1 are exhausted, the machine will be immediately set-up for type 2 production, even though the eventual decision may be to keep the machine idle. Now, when the machine is being set-up for item 1 production, no rewards are earned for a (random) duration d_1 . Then, production begins and kits for item 1 are exhausted. When the machine is set-up for item 2 production, no rewards will be earned again for a (random) duration d_2 . Hence, by switching to production of item 1, exhausting all kits of type 1, and switching again to item 2 production, the machine will have spent a total expected amount of time $D_1 + D_2 + \frac{x_1 + \lambda_1 D_1}{\mu_1 - \lambda_1}$, where the last term is the expected amount of time that the machine will be spending processing units of type 1. On average, however, rewards will be earned only for the expected duration of time equal to $\frac{x_1 + \lambda_1 D_1}{\mu_1 - \lambda_1}$. Therefore,

we define the reward rate of switching to production of part 1, exhausting the kits for part 1 production, and setting up for production of part 2 again as

$$\varphi = c_1\mu_1 \frac{\frac{x_1 + \lambda_1 D_1}{\mu_1 - \lambda_1}}{D_1 + \frac{x_1 + \lambda_1 D_1}{\mu_1 - \lambda_1} + D_2}.$$

Our heuristic allows switching to production of item 1 only when the index φ is high enough that, in the time interval where the machine is set-up for item 1, kits for item 1 are exhausted, and the machine is again set-up for item 2 production, the proportion of time that the machine is doing useful work (i.e., actually processing jobs and not being set-up) is at least ρ . In fact, it is easy to check that requiring $\varphi > \rho c_1\mu_1 + (1 - \rho)c_2\mu_2$ ensures this. Furthermore, Van Oyen and Duenyas (1992) show that this condition ensures that the system will be stable.

The derivation of the rule on whether to idle or to set-up when kits are exhausted, is based on a similiar idea. Assume that the server is set-up to process item 2 and $x_2 = 0$. In this case, we do not necessarily require that the machine perform “useful” work at least ρ fraction of the time when it switches, since the machine is not performing “useful” work by remaining idle. An immediate switch will still result in a “reward rate” of φ . However, consider the policy where the machine idles until the next arrival of a kit of type 1 and then is set-up for item 1 production. Of course, before the next arrival of a type 1 kit, type 2 kits might arrive, and rewards may be earned by processing units of item 2. However, we assume no rewards are earned while waiting for another kit of type 1, and compute the reward rate of the inadmissible policy that idles until the next arrival of a type 1 item, then switches to production of item 1, exhausts all type 1 kits, and again switches to item 2 production. This results in the following reward rate:

$$\varphi' = c_1\mu_1 \frac{\frac{x_1 + 1 + \lambda_1 ED_1}{\mu_1 - \lambda_1}}{\frac{x_1 + 1 + \lambda_1 ED_1}{\mu_1 - \lambda_1} + \frac{1}{\lambda_1} + ED_1 + ED_2}. \quad (2)$$

The condition for switching from 2 to 1 when $x_2 = 0$ is then given by

$$\varphi' < \varphi \quad (3)$$

which implies that the server earns rewards at a higher rate by switching now than by waiting for one more arrival of a kit of type 1. Simplifying (3) leads to a very simple formula for the

number of kits of type 1 required so that the machine will switch to type 1 production from type 2 production without idling:

$$x_1 > \lambda_1 D_2. \quad (4)$$

References

- Buyukkoc, C., Varaiya P., and Walrand J., (1985) The $c\mu$ -rule revisited, *Advances in Applied Probability*, **17**, 237-238.
- Hofri, M. and Ross, K.W., (1987) On the optimal control of two queues with server set-up times and its analysis, *SIAM Journal of Computing*, **16**, 399-420.
- Liu, Z., Nain, P., and Towsley, D. (1992) On optimal polling policies, to appear in *Queueing Systems*.
- Sarkar, D., and Zangwill, W.I. (1991) "Variance effects on cyclic production systems," *Management Science*, **37**, 444-453.
- Takagi, H., (1986) *Analysis of Polling Systems*, MIT Press, Cambridge, MA.
- Van Oyen, M.P., and Duenyas, I., "Control of a Queueing System with Sequence Dependent Set-Ups," Technical Report 92-60, Department of Industrial and Operations Engineering, The University of Michigan, Ann Arbor, MI 48109 (October 1992) submitted to *Management Science*.
- Zangwill, W.I., (1992) "The limits of Japanese Production Theory," *Interfaces*, **22** No:5, 14-25.