The Office for Control Over Results

An introduction to the cybernetics and practicalities of large scale modern production

By Donal Costello

black mold

The Office for Control Over Results is a blog written by Donal Costello, an Irish process control director for a factory producing electronic components. Donal has allowed us to reprint several of his essays from his blog to give our readers an introduction to what we might call progressive systems theory or cybernetics. Much of the socialist project has always been based on an abstract reworking of social labor, i.e. the way in which our labor relates to the broader reproduction of society. However, beginning in the latter half of the 20th century, new ways of conceiving of the concrete labor process and its relationship to self-management began to be theorized, perhaps most notably for our purposes by cybernetician Stafford Beer. The questions posed by Beer and others like him were from where does a theory of autonomy originate and how will this change our actual experience of work? For Beer, the answer to the first question was found by studying the variety, or measure of complexity, of a system:

> Senior Management is in principle unable to entertain the variety generated by any one (never mind all) of its subsidiary viable systems…The beginnings of a theory of autonomy, of decentralization, lie in this simple fact - rather than in political theory.

I think we here at black mold are adult enough to claim that the beginnings of a theory of autonomy are very much found both in systems and political theory (as well as a theory of morality), but Beer's point is a necessary one. When we talk about the need for a revolutionary transformation of society, what we are referring to is the fact that the productive forces of society have progressed to such an extent that the productive relations of society have begun to hold their further progression back. This is becoming nowhere more evident than in the capitalist organization of the factory, where outdated methods of hierarchical management systems based on the exploitation of labor (ie, class) are holding back efficiency and, more importantly, emancipated, creative and autonomous labor. Donal will explain why in the following essays. We will have a further reading section at the end for those of you who find that this short glimpse into the world of factory design and systems has piqued your interest.

-black mold editorial team

WELCOME TO THE OCR

My name is Donal Costello, I am from Dublin, Ireland. I hold an MSc in Engineering Cybernetics.

I primarily work on the design and implementation of lean production systems, inspired, among others by Stafford Beer and his management cybernetics and by the Toyota Production System (TPS) and its derivatives.

I am currently employed in this capacity by a large multinational manufacturer of domestic electrical components (plugs, sockets, switches of various types, etc).

My intention is to use this blog for three purposes: as a place to gather my own general thoughts, as a place where others can see the kind of work that I personally do, and most importantly, as a place where people can learn about the practicalities of large scale modern production.

Discussion of topics with the following headings is, as of now, my general action plan for this blog:

1. (Standard) time is money
2. The VSM – principles of homeostatic control
3. Why "control over results"?
4. The Toyota Production System (TPS)
5. Implications of modern technology and methods

If you're interested in a clear discussion of how technology, social relationships between individuals and groups, and 'results control' come together to regulate complex processes involving thousands of people – stick around.

(STANDARD) TIME IS MONEY

We're going to have one post on the economics of production – it won't be boring, and it's essential to understand the context of what automation and process control engineers are trying to achieve.

Let's say you take an economics course – you're going to hear a certain perspective on how the economy works. This perspective emphasizes the moment of sale. You have a product or service to sell (it doesn't matter how – maybe you bought it from someone else, or it fell from the back of a truck!), and bring it to 'the market'. Then the magic moment happens – someone buys it. They have your product, and you have their money. Supply meeting demand is the central fact of the economic process, and everything else is pretty much incidental to this.

The real story of how the economy works is very different. It's not that establishing the nature of in-demand products is unimportant (it is of course essential), but rather that this moment of sale is, from the perspective of a real producer, just a point on a continuum, which is neither more or less important than every other moment. By the end of this piece, I hope this will be clear.

Examination of a big company makes this point very obvious, because unlike a small company, it has clearly defined departments, which specialise in the different aspects of its operations. If you look at the organisational chart of such an enterprise, you will find a whole bunch of sections. Some of these are related to buying and selling – like supply chain management, procurement and customer satisfaction. Some are focused on production and the products, like research and development, process control and automation, and then you've got financial and management accounting, and finally auxiliary departments like real estate management, human resources, the legal office, and so on.

The company basically performs the role of a social-technical-economic control system, with one overriding aim in mind. This control system is operationalised, at the highest level of recursion, through financial accounting. The primary function of financial accounting is measuring the monetary flows of output, a 'profit and loss account', and the stock of capital (a 'balance sheet'), at a moment in time, to enable the owners, the shareholders, to calculate how profitable the company is, and to hold management accountable for this.

The same financial accounting principles and practices apply to all companies. They all must use them in order to obtain investment from the rest of the economy (the use of a standardised, objective metric like this for systematic regulation of the operation of a system is called 'results control' – which will be a common theme on this blog). An enterprise's Return on Investment (ROI) is the metric in this case – it is how investors and creditors hold management accountable for the capital and profit/loss positions of the firm.

All companies are held to account using this kind of results control, because the kinds of market portfolios which investors use are weighted shares of many companies. Investors and creditors demand a certain rate of return, adjusted for risk and other 'compensating' factors. Financial accountants know this as the 'required return'. This rate of return is consciously calculated in advance to reflect expected differences in rates of profitability over different time-frames, for different kinds of businesses. The econ-101 idealisation of bringing products to the market and hoping that they sell doesn't enter investor's minds at this level of abstraction – the big picture is worked out well in advance, with the risk-compensating adjustments built in.

These 'required returns' pose a practical problem for the boards of directors of companies – somehow you have to ensure that they are met. Those directors solve this problem by imposing (accounting) results control in turn onto their own firm. Management does this with the help of management accountants, who need to make clear the position of the company internally in terms of its ability to deliver profitability.

I mentioned earlier that the financial PL and balance sheets are calculated for a moment in time. The internal accounting results-control by contrast, as we will see, need to be continuous. In the same sense that if we want to drive to a destination, our 'internal' control over the car speed and direction must be continuous for the whole journey in order to get there.

Internal firm control system

Now another practical problem arises, and this is where it gets really interesting, because in order to realise the profit that the company is tasked with generating at the end of the process, the accountants need to describe the production process internal to the company in that same unit of account that the profit is described in – money. This is despite the fact that money plays no role within the company itself. Different departments internal to the company do not sell services to each other, and the workers and managers on production lines do not deal with money during their working day – they deal with time and quantities of product.

The depreciation of fixed assets can be dealt with simply in terms of replacement cost accounting, where asset values are described in terms of what it would currently cost to buy their equivalent and replace them. This gives you an objective basis for comparing the maintenance cost of the machine to the replacement cost, and knowing if you're on track, in this respect, to achieving your required return. But how can the use and waste of materials, and the actual labour processes within the firm be objectively taken into account when figuring out if the company is on track to get the appropriate rate of return?

Management accountants do this by distinguishing between actual costs in the enterprise, and what are called 'standard' costs. Standard cost expresses the use of resources, and labour time (of average skill and intensity of work, at the level of labour productivity called 'standard time' – more on this in a moment) in terms of money. Any real costs over standard costs are treated as a loss. Standard cost is therefore a measure of what the labour process, carried on within the company, should cost, if expressed in terms of money. This controls for the variability in actual day-to-day labour productivity and sets the benchmark for what would need to happen in order for the company to be profitable. Labour time in the firm is in this way reduced to the same cost unit as materials and capital goods (as accountants call it – 'absorption costing').

Because the accountants also track the value which is added to the product during production, the profit which they are charged with quantifying can likewise be found. Profit, strange as it may seem, exists during the entire production process, as a result of tracking the revenue and the costs, within the accounts of the company. Those categories do not just appear out of thin air at the moment in time when financial accounts are issued, detailing money income and expenditure. They are accounted the whole way through the process – which is why management accounting is essential. This means, in effect, that businesses keep two sets of books – one for the outside world and one for themselves.

Engineers with stopwatches

The automation and process control engineers build up a picture of 'standard time'. This includes several sub-categories – two of the most fundamental of which are 'useful time' and 'design time'. Useful time is that portion of the labour time during which work is actually being carried out on the product. 'Revenue' is essentially being created, just in the form of product now, not yet money – but this product is implicitly understood to be worth what our accountants and economists have projected in advance that we will sell it for – a realistic 'going market price' – so we speak of 'UT' as 'value adding labour'. Design time is the whole production process including not only work being directly carried out on the product but necessary work that does not directly produce products (like quality inspection). These are additional but unavoidable costs, not adding 'revenue'.

When the enterprise feels under financial pressure to hit its required ROI, and needs to reduce costs in existing production, or when it wants to start a new production process, engineers go to work. We use detailed studies of the necessary technical and labour inputs, design and 'value engineering' studies, observation based on trial runs (as a last resort with video cameras and stopwatches), and work studies. The gold standard is MTM data (which are kind of standardised building blocks of how long it takes an average human to do all kinds of manual tasks). We design and redesign the production processes on the basis of concrete but ideal plans, and then try to reach those ideals in reality (but we never can of course, because there are many kinds of time-wasting that can be reduced in many ways but never eliminated).

By multiplying these standard time measurements by the wage costs of the relevant workers, we arrive at the standard cost of the labour inputs used to make the products – what the management accountants are interested in. So while it is widely understood that process control and automation engineers design production processes, what is less widely known is how this fits into the bigger picture. We organise the technical control system, which fits into both an accounting and financial control system – the results of which are all expressed in terms of money, so we can track these money categories all the way through.

A single system

Companies do of course nervously watch whether their products sell in the market or not – and sometimes make huge blunders in what they produce, or sometimes just get unlucky (think of an event like Covid 19 and what it did to many industries at that time). But companies spend just as much time on their internal processes, and for good reason: The commodities which we make are expressed in terms of their prices before they ever get to the market, as a pre-planned precondition of the process of which sale for money is just one moment. Profit is similarly produced before the outputs go to the market. The form of the profit just subsequently changes from product to money (and later, when reinvested, back into product again). This profit, made in production, can be fully or partially lost during or after production.

So yes, it is of course true that if finished goods are unable to be sold, the company has lost money. Costs outstripping sales is what a 'loss' means in terms of accounting profit. But the same phenomenon is recorded at every point in the production and circulation process, sometimes as accounting profit and sometimes as profit in internal management accounting reports. If products are destroyed during manufacture, it's also a loss in the same sense as the unsold product. If on a production line, we experience 'unscheduled downtime' (the arch-enemy of process control engineers!), we have deviated from our standard time by that much, and this is, in exactly the same sense, a loss – we are paying real, measured money costs for minutes of production during which products were not produced. This is a cost above the standard cost without a corresponding flow of revenue. Or, suppose we successfully sell our products for the full intended money profit, and as a result have a nice big bag of cash, but then this cash is subsequently burned up in a fire before it can be reinvested in new production. This is a loss to the company in the same sense as in the other instances. All are points on a continuum – none are more or less significant than the others.

This blog is about control over the technical and labour process, and now we have put that control system in its proper context. I hope it made you think!

THE VSM - PRINCIPLES OF HOMEOSTATIC CONTROL

Any YouTube video that you watch on the topic of cybernetics will inform you that the term itself comes from an ancient Greek word for the 'helmsman' of a ship, and that it is a multidisciplinary field focusing on "homeostasis" and "autopoiesis". In other words, in a nutshell cybernetics is the study of systems which modify their own behavior in order to continue to operate viably in their environment, and more than this, that when they don't know how to suitably modify their behavior, they self-organise a discovery process to figure it out.

This, as you can imagine, casts a wide net over a lot of kinds of systems which could potentially exhibit this behavior to some extent. Some of the most obvious ones are biological systems like ourselves. But technical (mechanical,electronic,informational etc) systems also have strong cybernetic elements today, through self-diagnosis and self-correction subsystems, although not of course at the level of sophistication found in biology. This kind of 'engineering cybernetics' is what I myself studied in university.

But a different focal point in cybernetics is our topic of interest here – that being management cybernetics, the explicit recognition of organisations as socio-technical systems conforming to the same rules which apply in principle to all cybernetic systems. Stafford Beer was, in the western world, the great pioneer of management cybernetics (I say in the 'western' world, because others were developing similar ideas, which had a greater overall influence, at the same time in Japan, and we'll consider that in a future piece).

Beer's great work was the Viable System Model (VSM). The VSM is still quite influential today, and this is a testament to Beer's achievement. I can personally attest to the fact that Fortune 500 companies are using the Viable System Model in its unaltered form, because there is simply no better way of cutting through the pathologies which exist in any organisation and showing the true state of how all of the essential parts of a production system actually work.

Wikipedia has a reasonably good and brief explainer for the VSM. My even more condensed explanation for our purposes is that any viable system, be it the human body, a social organisation or anything else that exhibits cybernetic properties, can be represented by the VSM. In other words, the VSM is a visual mapping of the elements which every cybernetic system has (rather than an idealisation of what the system should have). If it is a functional cybernetic system then it does in fact have these components, although they may not be organised optimally.

Also, as a visualisation, the VSM only captures a specific system in focus. In management cybernetics therefore you might just be interested in looking at one company as a viable system (and it will be the 'system in focus'), or a whole town, or an industrial sector, etc.

The VSM shows the active elements of the system, along with the regulators of those active elements. If your system in focus is a company, and within that company there are offices in which work is being done, these offices are the active elements – called the 'system 1s'. The lower management functions (systems 2,3) are the automatic regulators for the system 1s. For example if the accounting department notices an error in their work, they have a local process to go about resolving it. If a computer isn't working correctly, there is an IT specialist who can remotely fix it for them. There's no need for higher management functions to enter into the equation here.

Higher management functions in principle do not interfere in, and are not even aware of, lower management functions' specific activities, unless a certain predefined threshold of importance is reached, and then they absolutely must become involved.

These higher management functions (systems 4,5) do a lot of the conscious, deliberate, not strictly reactive work of the system. Things like conscious, non-routine interactions with the external (in this case, business) environment, planning the direction of the system and establishing global goals and criteria for the whole system. They receive only the most important, highly filtered information from the lower levels, and use this information to make executive decisions, including the normalisation of any problems which arise in the system when its lower management functions are unable to resolve them.

So basically, you can think of the VSM as a set of roles, or functions. These are roles that must be performed by the organisation, but there is nothing fixed in stone about how exactly they are organised – which is a feature, not a bug. In management cybernetics, homeostasis is sought as the ideal type of control. Here we are interested in self-correcting behaviour and the self-organising of solutions to maintain stability in the face of a changing enviornment. This type of control is regarded as infinitely superior to top-down control, because it involves the entire system in the problem-solving effort, and thus there is no sense in which the higher management functions are more important than the lower management functions and active elements. Each role is equally critical, and each shapes the development of the system organically – something that can seem quite unsettling for some managers who would like to imagine only themselves as defining how the organisation works.

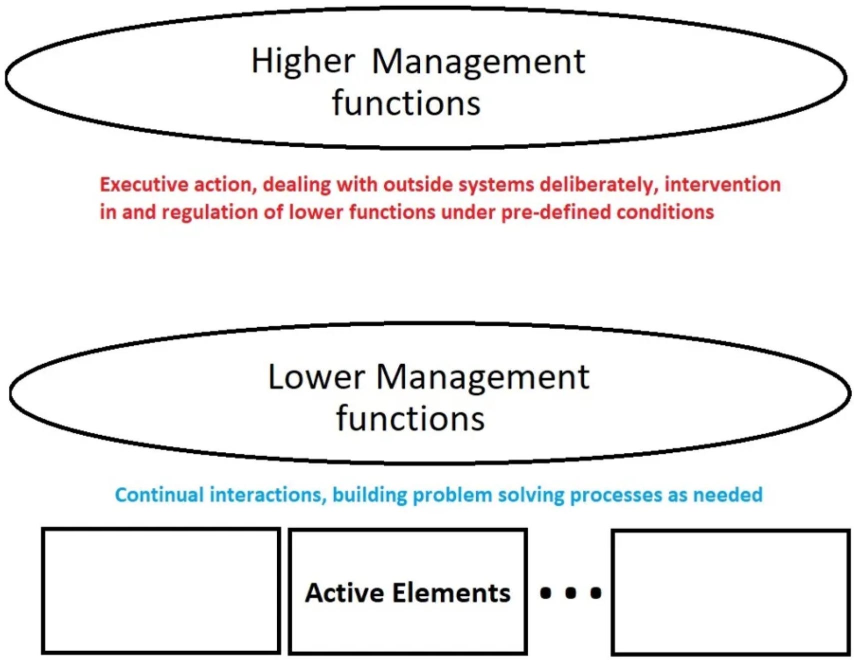

To wrap up our introductory explainer, we need to emphasise the issue of 'requisite variety'. A company using a homeostatic control system as we have briefly discussed must ensure that only the right information is going where it needs to go, because the bandwidth of any element of the system is quite limited. Higher and lower management functions must not be cognisant of all information happening within the system 1s, but only that which rises to the level of a problem (or initiative) suitable to their level of responsibility for normalisation (or implementation). So one of the very important elements of designing a socio-technical control system for your organisation is the clear demarcation of under which circumstances different elements of the system must intervene and assume responsibility, as this defines what information will be allowed to pass through the filters between management elements and reach higher levels.

WHY "CONTROL OVER RESULTS"?

I realise that describing my blog as being concerned with "control over results" may seem like a pretentious way of saying that it concerns "management". So I want to promise you now at the outset that no, this blog is not a conspiracy to waste your time. What is almost always meant when organisation speaks about "management" or "to manage" something is a catastrophic concept that I wouldn't want to be associated with.

If you have not read my previous post on the VSM, I want to reiterate that in any organisation:

> System 1s are the active elements carrying out the real purpose of the organisation (eg an assembly line in a factory).
>
> Lower Management functions are concerned with continual interactions with system 1s, building problem solving processes as needed locally.
>
> Higher Management functions are concerned with executive action, dealing with outside systems consciously/deliberately, intervention in and regulation of lower management functions under pre-defined conditions.

Cybernetic 'results control' is distinguished from 'bureaucratic control' in the following sense.

Bureaucratic control is control over an organization by processes to which the decisions of higher management functions are the 'inputs' to the control system. A bureaucracy, seen in this light, is nothing other than the higher management functions controlling the active elements (system 1s) directly. Special groups for this purpose deal only with the variety of situations that are prescribed by the higher management functions, and handle them in a pre-determined way, with each special department dealing with a particular class of issue. Issues that fall outside of those responsibilities are not catered for by the organisation at all. Bureaucratic control is therefore a static and discrete-time controller, established and perhaps periodically modified somewhat, acting out the instructions of the organisation's higher management functions.

Results control on the other hand can be thought of as control over the organization by processes to which (especially) system 1 process results are the 'inputs' to the control system. This is a fundamental aspect of cybernetic (homeostasis-seeking) control. Results control seeks a situation in which any amount of variety can be handled. It is a dynamic and continuous controller, because whatever issues emerge in the system 1s, it leads to someone having managerial responsibility to create a process/solution for this issue. The role of management functions (higher and lower) in a cybernetic / results-oriented control system reduces to (1) definition of the conditions under which the given managerial functions become actively involved, and (2) specific actions to take in each type of case where they have had to become involved.

While in results control, many management responses will become standardised in the main categories of issues which continually arise, this is altogether different to the higher management functions prescribing in advance what kinds of issues the system 1s are permitted to deal with, as in bureaucratic management.

Let's consider the example of a production line in a factory, to clarify what all this means in practice.

A bureaucratic control system applied here would be when the higher management functions of the factory view the improvement of the productivity of the system 1 element here (the production line) as a technical problem which they will solve. A special group of people will be convened by the given higher management function, and they will figure out ways, within pre-defined limits, of increasing the work intensity on the line, for example though financial inducements and punishments. The measurement tools which are often used here are physical output, working time, survey data and so on. Again it is a static set of arrangements which acts directly on behalf of higher management.

A cybernetic oriented system would work very differently. Firstly, the control system is based on process results, so you have to choose what you measure in order to gain more insight into how this production line can operate without the deeply unhelpful involvement of bureaucracies. Suppose that a system exists whereby the operators on this line can declare whenever they finish a batch of products they made, and in what window of time they did this work. Allowing only for this, consider some simple measures that we can now use:

1. Operator Time (OT): This is just the amount of time that the operators on the production line described themselves as working on the completion of a given quantity of product. Typically therefore it might be 'from 9am until a break at 11am', and the output might be 100 units of product. If there are 3 operators on the production line, then OT in this case will be 2 hours × 3 = 6 hours.
2. Actual Production (AP): The quantity of produced product during the OT, 100 units in this case.
3. Design Time (DT): The amount of time in which, using the current production methods, under normal production conditions, assuming an averagely skilled operator working at a reasonably relaxed pace, that the production process can finish 1 unit of product. (In reality this is made using a weighted average, because there is not always a linear relationship between time and output quantity).
4. Standard Time (ST): AP × DT. The amount of time it should have taken to make the 'Actual Production' quantity of product, given DT.
5. Plan Production Quantity (PPQ): OT / DT. The amount of product that should have been made, given the operator time (OT) on the production line.
6. Efficiency Coefficient (EC): ST / OT. The relationship between the operator time on the production line, and the amount of time that should have been spent on the production line, given the production result.

When the operators on the production line record a new production batch of some number of units as having been completed, the quantity that should have been made in that time is already known (PPQ).

This can be compared directly to the factual quantity that was made (AP).

Simultaneously, we learn the amount of time that it should have taken to produce the production result and also the factual amount of time that was taken (ST and OT respectively). These can be compared directly, which is called the Efficiency Coefficient (EC). It is advantageous to an operator to clock up more OT (as salaries are generally paid for the amount of clocked-in hours), but this higher OT requires correspondingly higher output during the working time, or else the EC will be seen to fall.

To the extent that OT > ST, this is known as "red time". Red time begins life as "unclassified red time", a given amount of time that for some reason did not result in the production outcome that we had hoped for. Classifying red time means explaining it qualitatively (e.g. "we were waiting for parts to arrive from the warehouse"). This act of explaining is another measurable activity which now becomes one of the primary concerns of the lower management functions in our factory. This function could be represented by a technician who simply comes by the production line and helps out whenever she/he notices that red time for this line is above a certain threshold.

By helping out and understanding the problem, the technician is in a position to explain it, something which the operators themselves can not be distracted with. We have thus gained a new parameter: "% of classified red time", derived from the normal activities of our technician. As portions of unclassified red time are selected and explained by the technician on her/his computer, this percentage indicator and the types of 'red time problems' that arise now become the key to 'results control' for the technician's process, and become the basis for intervention by other management functions under predefined criteria.

We can note that we have assembled a quite comprehensive picture of the results of this production process, and we have done so without a bureaucracy having to be established. Simply by engaging in the action of producing something and logging it, our database has been populated so that both lower and higher management functions now face a much simpler task. They just need to decide under what circumstances they would have to become involved in the life of this production line, and, given different scenarios, how they might proceed from there.

The reason that 'results control' is so advantageous compared to a bureaucracy is that any process result creates the need for a response that can then be defined by the lower and higher management functions. The results on the production line might range from 100% absenteeism, to real production being 300% higher than anticipated – but whatever the actual event, it can be handled and does not require a bureaucracy to delimit the possibilities beforehand. That being the case, continuous and dynamic self management becomes possible.

But we must note something very important in this case, which is indicative of cybernetic management in general. DT (the amount of time it should take to make 1 unit of the product) is not a goal or target value. It is an understanding of the amount of time required to do the job comfortably by the average person, within the context of the task. Therefore higher productivity is not demanded if operators consistently exceed the DT productivity – as any bureaucracy appointed by a higher management function would immediately jump to do.

The productivity level represented by the DT is that which allows for the self-development of the process in a somewhat relaxed manner. This is measured by the Efficiency Coefficient (EC), which needs only to be kept in check. Unless the process itself is fundamentally changed by production engineering experts, the DT does not change, and so workers never need to worry that the Design Time will decrease if they find a flow of work that allows them to consistently produce faster than it would suggest.

The logic which holds true for a single production line also does for the organisation as a whole. Measures such as the ratio between the number of non-productive work (NPW) hours and the hours worked in direct production provides the basis of control over the results for the whole organisation. This kind of high level indicator is a good basis for higher management functions to intervene when necessary, if the proportions cross certain thresholds.

Consider one final example to appreciate the intuition behind cybernetic results control. In our modern societies today, we have bureaucratically maintained health services, which are comprised of departments with which sick patients should make appointments, within predefined categories. They also maintain additional bureaucracies established by higher management functions (both government and private) to bombard the health professionals with paperwork and try to assess their performance.

A results control orientation might instead rely on conducting continual and general medical checks of the population, the results of which are the inputs to the treatment scheduling process for more specific tests and treatment. Now there is no duplication of effort or activity – the process itself is what creates the performance indicators, so that doctors, nurses and so on never need to consider that a process apart from treating patients is even happening. They just log basic information and when times are difficult, they notice that management functions which otherwise leave them alone take a sudden interest (needing to, as their processes are also subject to control), and that increased resources and help then duly arrive.

It is simply a non-obvious conclusion that the best outcomes by far come from empowering people through establishing results control over suitable results parameters, and allowing processes to develop which best 'normalise' those parameters, rather than decreeing static processes from the top down. This is why people hate dealing with organisations in general, because by killing variety at the system 1 level, (rather than allowing the system 1s to be as flexible as possible and then reducing variety through subsequent information filtering), an organisation becomes intolerable to deal with and quite useless for those both inside and outside of it.

Unfortunately but understandably, people in higher management functions almost always believe that they intuitively know best how to handle every situation in 'their' organisation. This being the case, they establish these static bureaucracies which are charged with dealing with specific situations in specific ways, as proxy actors. Don't do this – instead, step into the office of control over results!

In the next post I will be discussing the Toyota Production System (TPS) – a great leap in cybernetic management theory, where the principles of homeostatic results control are taken all the way to their logical conclusions, as a planning tool which allows the organisation in many important ways to run itself and continually improve its performance.

LESSONS FROM THE TOYOTA PRODUCTION SYSTEM

The Toyota Production System (TPS) is about as far as our human industrial civilization has so far advanced in creating a cybernetic management system for production processes. Whereas Stafford Beer demonstrated which functions an optimal cybernetic system needed, the TPS actually built out the functions for their specific economic-technical and social context.

The TPS emerged from the need which Japanese industry faced to catch up with the huge US automobile manufacturers from the 1950s onward. During their tours of the United States, representatives of the Toyota company noticed two important things – firstly that the Fordist and Taylorist manufacturing orthodoxy looked good on the surface but was actually extremely wasteful, and secondly that the key to a superior system lay in how US supermarkets implemented stock control in order to provide the right amount of product to meet customer demand.

It was not a coincidence that this happened in the automobile industry. It was here that the complexity and variety within the production processes reached such a high level, that in order to make significant further improvements, a new management system was objectively required. Toyota discovered, rather than invented, this path forward for high complexity production – the Fordist technique, as we'll now consider, had already been played out.

What was the problem with Fordism?

Simply a lack of variety. Fordist production had, from 1908 onward, represented a massive jump in quality and productivity over how factories used to operate, essentially as large workshops. The idea of training people into specialised roles, using standardised tools, to produce exactly the same product over and over on an assembly line, allowed for the first time products of fairly high complexity to be created on a mass scale with a reproducible level of quality.

Workers were regarded here as a type of machine, which could produce a certain number of output parts, given the presence of other machines and a certain number of input parts. This was spelled out explicitly in Frederick Taylor's Principles of Scientific Management (1911). Production engineers would pre-determine every minute detail of what every worker should do, and then the workers would follow those work steps precisely. That was very successful, and is still a widely used method today, especially in low-technology production processes.

However, there are major drawbacks. For example, any time you want to 'improve' a process in Fordist production, it involves industrial engineers undertaking a long period of designing, trialling and modifying, and then training workers and standardising the new method – this is difficult and time consuming. Faced with the pressures of expanding production rather than tinkering with and disrupting it, management will try to minimise changes to the existing production system. This is the opposite to what you would like in a cybernetic system, which should continually and reactively self-improve and adapt its operations, rather than improvements only coming in a top-down way as a result of proactive management decisions to change processes.

So this was the first big drawback with Fordist production – it requires active management by (in cybernetic terms, which we covered in the piece on the VSM) the 'higher management' functions of the organisation. At a certain level of complexity for any cybernetic system, this becomes ineffective. Applied to the human body, this would be like the higher management element (your conscious mind) directly trying to control the rate at which your heart beats, lungs breath and eyes blink, rather than having you body's lower management functions 'reactively' self-regulate those mechanisms.

I have mentioned that Fordist production is also wasteful. The reason for this is pretty simple. Capacity utilisation is generally how the real production situation is compared to a pre-planned production outcome. If capacity utilisation is very high, for every machine and production line, it means the system is working well, and the plan targets will be achieved or even surpassed. The problem with that kind of proactive planning is that, because each process has a different production rate per unit, and each process is trying to operate at maximum capacity, it means that slower production processes will quickly be inundated with input parts that they can't quickly clear. Faster processes will have the opposite problem – they'll be waiting around for inputs to be ready, and so to protect against this, they try to build up large inventory buffers.

This means that large stockpiles of 'work in progress' (WIP) inventories exist throughout the supply chain, both for slow and fast process steps! Hidden among those stockpiles will be defective products, and if you find one, it is advantageous to simply throw away the defective component and move on, rather than delay production by investigating the cause of the defect. You will also end up with a lot of obsolete stockpiles that you just need to dump.

But most problematic is that emergent problems in the production process do not dynamically force a change in how production works, because there are always plenty of spare components to keep the process going at a high utilisation rate. Even if a part of the process fails for a whole day, this is unlikely to force a change in how the process works, because that would mean a lot of extra, and apparently unnecessary, work. On paper, it doesn't appear urgent to change anything because if you look at the production data, you'll see that thanks to inventory buffers, production elsewhere never stopped during a period of unscheduled downtime at any given step in the process.

What I have described above is these days called 'push' or 'batch-and-queue' production, but is synonymous with the old Fordist system. It is called 'push' because each process pushes its output onto the next step as fast as it can, causing a high utilisation rate and low cost-per-unit, but also lots of wasted time, effort and material, and a structurally built-in resistance to change. It is also the 'natural' system from the perspective of higher management control, because they assume direct responsibility over each stage in the production process.

Production from below

The Toyota Production System is fundamentally a self-organising system, and therefore a truly cybernetic organisation of production. In contrast to Taylorism and Fordism, in which production engineers design the minutia of specific production operations which are then carried out to the letter by factory operators, here the operators themselves 'hold the stopwatch'. Workers on the factory floor are organised into 'work teams', with a 'team leader' (who is more like a 'delegate' than a 'leader'). The leader's job is to gain consensus with their team about the best current way to carry through production. The leader is responsible for it, but does not have the authority to keep their new method, if the higher management function of the organisation decides that this method is not appropriate for any reason. The team leader has responsibility, but not ultimate authority.

This means that the production process under the TPS is constantly changing, not according to the direct instructions of the factory's higher management but rather under their close supervision. The role of the process control engineer now reduces to interfacing with the work team leaders, and standardising the current best practices that they have come up with. On the basis of this process of moving from one standard to the next, while always improving, the TPS enables highly complex production systems to adapt as needed, superseding the rigidity of Fordist production.

It is also important to consider the social-relations aspect of this. The implication of a continually improving production process, from the perspective of the operator, is that they might easily 'improve themselves out of a job'! This is always a major practical problem with the implementation of the TPS – management think that they can get workers to improve processes, without worrying about the implications of this for their own employment. For the TPS to work properly, the workers must be given guarantees that they will not suffer any negative consequences as a result of their own process improvements. Most famously this is expressed through Toyota's 'no layoffs' policy, which they have at various junctures at least tried to implement.

Pull, don't push

The Toyota Production System is a ‘pull’ production system, again, modeled on supermarket stock control. This relies on Deming’s principle that ‘the next step is the customer’. Unlike in a supermarket, here the inventory buffers exist in-between production steps, so that the process takes on the form of a chain of steps, with a buffer separating each ‘supplier’ from its ‘customer’.

Specifically in what is called the ‘Kanban’ system, each buffer is a group of ‘bins’ of components, and each of these bins can contain only a fixed quantity of parts. So the idea is that the ‘supplier’ production step can only take away a bin and refill it, on the condition that the ‘customer’ production step has returned the bin empty. The time available to the supplier to produce the products (called the ‘takt time’) is in this way forced to synchronise with the rate of consumption of the customer.

In the event that the bins are replenished by the supplier at a slower rate than they were consumed by the customer, the latter would soon run out of input components and have to close down their production line. Each subsequent step would also have to close down given their lack of inputs, and each preceding step, due to their inability to offload their outputs. Equally, if parts are provided by the supplier more quickly than the customer consumed them, then the bins would quickly fill up, forcing the supplier to shut down its production line when all of the bins were already full. Sychronisation between all steps is thus unavoidable.

So unlike Fordist production, using the TPS we have no interest in maximising productivity at any given step in the process – we only care about total throughput, even if this means machinery lying idle. This implies that we should redesign the supply chain to achieve an ‘even’ rate of production at each step, which is equal to the production speed of the slowest step. We should then work on improving production at this step.

A brief side-note: Referring back to the production metrics in *my post called 'Why "control over results"?', my spin on this is that production should basically conform to the following equation. If it does so, then, during each 'takt time' in which the buffer is refilled, supply and demand are forced to match.*

(OT / DT) / (shift time / takt time) ≈ average demand / (shift time / takt time) ≈ buffer capacity

or, Supply ≈ Demand ≈ Buffer

[Where OT = 'Operator Time' (the total hours worked during a shift), DT = 'Design Time' (current standardised time to make 1 unit of the product), 'shift time' = number of hours in shift, 'takt time' = time taken to replenish buffer, 'average demand' = the current average level of withdrawals from the buffer during a shift]

Process improvement

Assuming that the buffer consists of a given number of inventory bins, Toyota quickly learned that the removal of one of these bins would force the rate of replenishment of each remaining bin to increase. This is clear in the above equation. If the buffer capacity is reduced, then, with demand and the shift time unchanged, the takt time must fall.

This meant that the margin for error in the synchronisation of the production steps was reduced. Now, one fewer bin could be restocked by the supplier, before it would be forced to shut down production due to insufficient demand. From the perspective of the 'customer' step, the time needed to consume one fewer bin full of products now represented the maximum amount of time that their production could continue, if supply was interrupted for any reason.

It was found that work teams are extremely innovative at redesigning production processes over and over, in order to keep production going under ever-tighter time constraints, established by this reduction in the size of buffers. Thus the supermarket 'pull' system came together with the 'self-organised' production philosophy. As a direct result, under the TPS the level of quality rises rapidly beyond that achievable under traditional Fordist production, because a faulty component is the same as no component at all. A batch of faulty parts would equally cause the shutdown of the subsequent production line, as these are observed by the 'customer' step right away who determine that they cannot be used (rather than added to a stockpile).

Algedonics
What Toyota called the 'Andon' signalling system is Stafford Beer's 'algedonic signal', discussed in my piece on the VSM. Algedonics are necessary for the stock-control and self-improvement system to function, as they are necessary for any good cybernetic system. An Andon signal alerts those to which it is addressed that a problem has arisen which requires their immediate attention.

Whenever it was found that the 'Kanban' bins were either all empty or all full for example, resulting in the stopping of production, an Andon signal would be activated. The signal would at first be addressed to the work teams immediately responsible for those lines. From the moment of its activation, a timer would begin, giving the work teams a period of time in which to resolve the problem. Doing so quickly was critical to preventing the whole supply chain from having to stop production. If they were unable to do so within a short period of minutes however, the Andon signal would 'escalate' to the next level of the organisation's hierarchy. Repair personnel and production experts would be alerted that they need to come and help resolve the problem. If they too were unable to resolve it within another fixed time period, it would 'escalate' to the next level of responsibility and so on.

The occurrence of the 'problem' itself was found to be the key to the self-improvement process. While a Fordist production system would try to avoid problems which might interrupt the production process, here the problem which had interrupted production becomes the basis for the next round of improvement to the process. Without it, the improvement would never have taken place. The 'Andon' system is thus a classic example of an 'algedonic loop' modifying the behaviour of a cybernetic system. This critical cybernetic feature is completely absent from Fordist production.

TPS Today

If you're reading this on any modern electronic device, then you've already got a product which was produced using the TPS or one of its derivative systems. In medium to high complexity production, Fordism and Taylorism as philosophies are no longer relevant, and the TPS has become ubiquitous. The factory where I work can be characterised as of 'medium' complexity, as we produce domestic-use electrical equipment, and for us the TPS is undoubtedly the standard to which we try to hold ourselves.

In lower complexity production, this is not generally so. Sweatshops making jeans in Bangladesh do not feel the need to deviate from pre-TPS production principles. I believe that due to the perception of a loss of control by managers, such radical changes occur only when the conditions are such that an industry has no choice but to adopt them. Toyota themselves, after all, from the beginning adopted such a radical strategy when they were facing bankruptcy in the 1950s, and by the 1980s all of their competitors, from Ford to GM, seeing the writing on the wall, copied the TPS as best they could.

There's vastly more to say on the TPS, but rather than go any further here, I encourage everyone who hasn't done so to pick up the classic texts on the subject, *The Toyota Way* and *Lean Thinking*.

Further Reading

Donal's blog can be found at theofficeforcontroloverresults.wordpress.com.

Additionally, there are several texts that we at black mold would recommend on the subject of systems theory and design. These are:

- Modern Industry and Prospects for Socialism by Donal Costello (Cosmonaut Magazine)
- The Volvo Uddevalla Plant by Christian Berggren
- Diagnosing the System for Organizations by Stafford Beer
- Designing Freedom by Stafford Beer
- The Dialectical Biologist by Lewontin and Levins

Donal Costello is a process control director from Dublin.

black mold is a group of like-minded anarchists, communists and other nogoodniks who seek out and publish political work as they see fit.

blackmoldpress@proton.me

blackmoldpress.com

The Office for Control Over Results is an introduction to the organizational systems that make up modern factory systems, as well as the potential for their restructuring into systems based in autonomy and self-improvement.

black mold