

Module 1: Foundations

A different way of thinking about 'resource allocation'

The challenge I have set myself is to redefine the way that we control resources – expenditure – in organisations. The way in which I (and I am sure) you were taught to do this was through the process of annual budgeting.

Budgeting is so deeply ingrained in the practices and thinking of those that manage finance that it is difficult for them to contemplate not having one. So much so that many people think it is the only way to exercise control, over costs or anything else involving money. So, the challenge is more than just about defining a new set of tools and methods. We need to think about 'control' in a different way. One way to reboot our brains is to look at it from a different perspective. Here goes.

My starting premise is that money plays the same role in economic organisations that energy does in organisms. Energy is defined as 'the capacity to do work' and I define economic organisations as those that have making money as a core (but not necessarily exclusive) purpose. So, this analogy seems to fit. The only way that work gets done in organisations is by paying for it. Let's look at how organisms control energy in order to do work.

Energy in the form of sunlight (if you are a plant) or food (if you are animal) is put to five types of uses:

1. To maintain the existing organisational structure
2. To change the structure of the organisation – its size or morphology
3. To acquire more energy
4. As storage for the individual organism (e.g, fat) or future generations (e.g, starch in seeds). And finally, for lucky organisms like ourselves...
5. To achieve another purpose...such as raising children, caring for the sick, dancing, writing books etc.

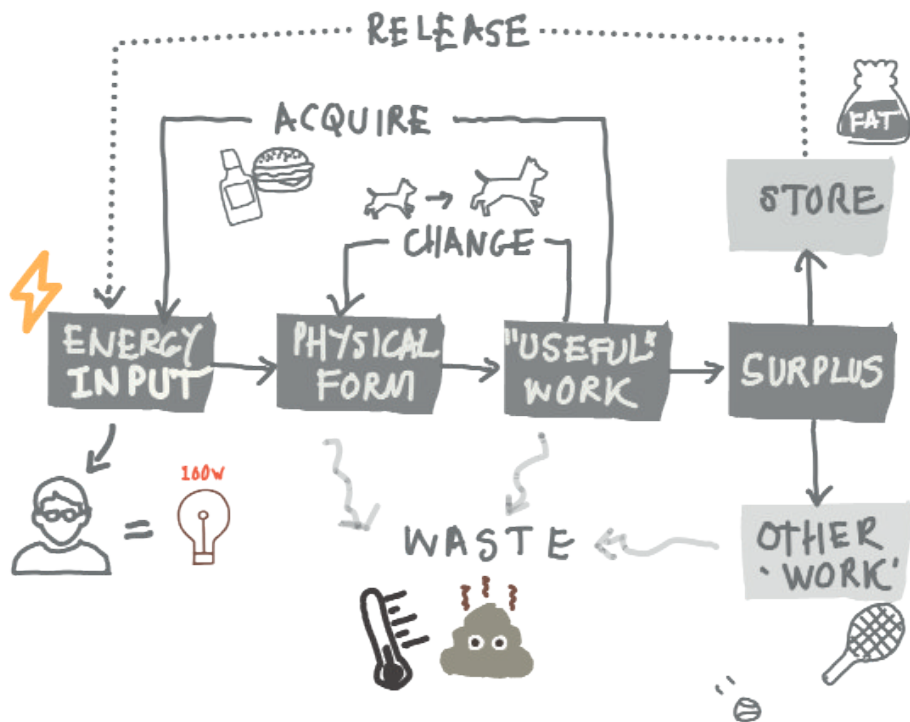
The allocation of resources to these different ends is controlled by the organisms metabolism, and every organism on the earth today have extraordinarily efficient metabolic processes, courtesy of 4 billion years of evolution. Energy that is not put to good use is waste. In organisms waste takes the form of low-grade heat, excretions and, in animals, unnecessary effort. So low levels of waste are the hallmark of an efficient resource management process.

But survival is more than about the efficiency of internal processes. To survive, organisms must change what they do and how they do it to compete effectively for resources in the environment.

The metabolic processes of organisms are good at allocating (energy) resources to preserve life and adapt to the environment, meaning that there is little waste.

Energy flows in organisms

Regulated by **metabolism**



$$\text{ENERGY EFFICIENCY} = \frac{\text{INPUT} - \text{WASTE}}{\text{INPUT}}$$

le. main tenance + 'useful' work

How does an organisation metabolise energy (money)?

Although this isn't how we describe them, organisations also have metabolic processes that direct the flow of resources with the aim of doing the right things at the right time and in the right (most efficient) way. These processes are usually called something like 'cost control' or 'resource allocation', but I prefer the term 'regulation' because it doesn't have the negative connotations of the word 'control'.

The energy (money) that an organisation needs to survive comes from providing goods and services to other organisations (customers or clients) that they are willing to pay for. This consumes resources, as does the requirement to maintain the structure of the organisation. Any surplus money can be used to add to or change the existing structure, stored for future use, or put to another purpose. This might involve repaying those that provided the resources needed to set up the organisation in the first place, such as banks or shareholders.

In an economic organisation waste is defined as anything that has no value to a customer (and so is not prepared to pay for) and is not necessary to maintain or change the (product or process) infrastructure of the organisation. If the level of waste is too high, or resources are misdirected and the organisation will either die (go bankrupt) or be consumed by a predator (taken over). Unfortunately, it is not always easy to identify and measure waste in economic organisations.

Over billions of years biological organisms have evolved clever ways of regulating energy flows. Fortunately, we can short circuit this process by learning. Our challenge is how best to maintain organisational viability (its capacity to survive in the short medium and long term) by regulating the amount of money:

1. **Consumed** by the value generating processes
2. Incurred in **maintaining** the existing structures
3. Spent on making **changes** to the above
4. Directed to **other purposes**.
5. Held in **reserve**

We need information to make intelligent regulatory decisions. What information do we currently have and how should it be used?

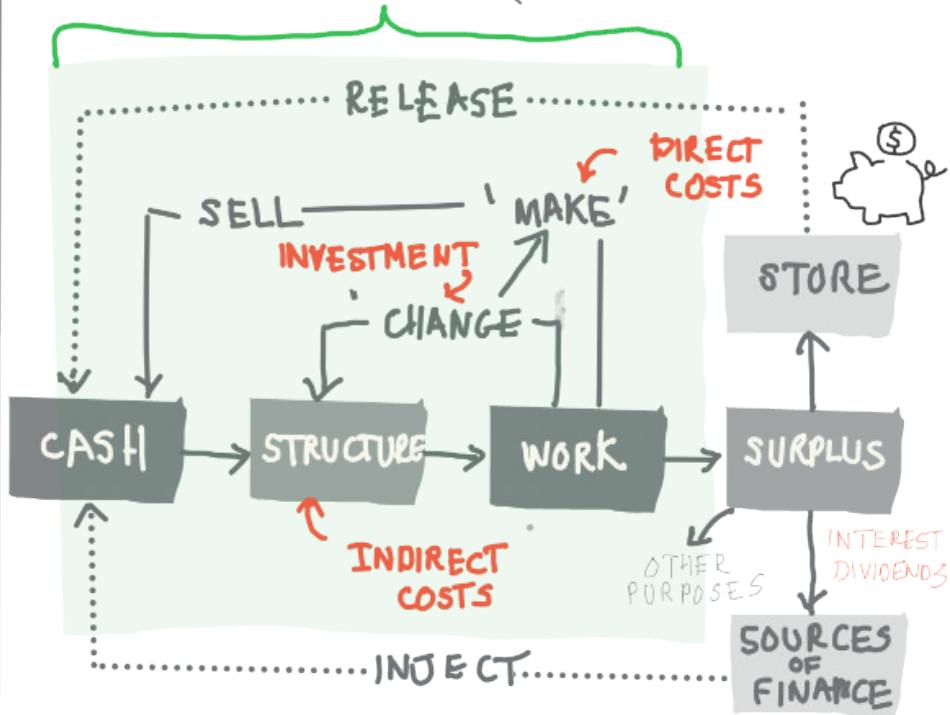
Most of what follows in this section will be very familiar to cost accountants, but to avoid confusion we need to ensure that we have solid foundations, based on a shared understanding and common definitions.

Regulatory mechanisms are required to direct how organisations consume energy (money).

Money flows in organisations

↓
Regulated by control processes

MY FOCUS 🔍



WASTE IS
DISORDER

How cost data is captured and categorised.

Before we can generate information that can be used to make decisions, we need to collect and categorise data. We need to know what, where and when resources have been consumed. This is a simple matter of record keeping involving allocating transactions to accounts (e.g., labour, materials), cost centres (e.g., production, sales head office) and time buckets (months, years). But to convert this raw data into information that can be used to regulate the consumption of resources we need to know more than this.

First, we need to know what we can **ATTRIBUTE** costs to. The most important distinction is between **DIRECT** costs, which are those that are directly attributable to the production of goods and services, and **INDIRECT** costs that are not.

Secondly, we need to know how costs **BEHAVE**, particularly when the volume of goods and services changes. Specifically, we need to know what proportion of direct costs can be classified as being **VARIABLE** with output. All other costs (including overheads) are called **FIXED** costs.

Within the fixed costs category an additional distinction can be made to reflect the degree of **CHOICE** we have. In most businesses, most costs are **NON DISCRETIONARY** since they are driven by how the business is structured and the nature of the goods or services it sells. But there will always be some costs incurred by choice. These are **DISCRETIONARY** costs.

So far this seems straight forward. But in practice things can get complicated. For example, different types of businesses can have very different cost profiles.

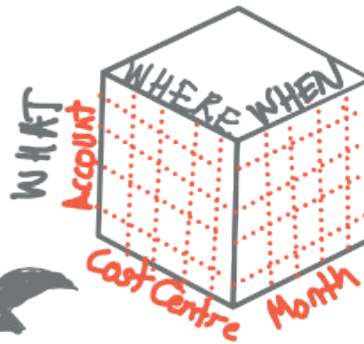
A high street retail business will tend to have a high level of direct costs, much of which is variable because they purchase everything that they sell. On the other hand, the variable cost of a software business might be close to zero. And the definition of what is direct and indirect, variable, and fixed is often ambiguous. In particular, what is treated as a variable cost depends on the frame of reference. For example, a cost categorised as 'fixed' in the short term can be variable over the long term or when there are large changes in volume.

Furthermore, the relationship between changes in costs and volume is rarely a simple linear one. Economies of scale in material procurement being often results in lower marginal costs as volume increases. And 'fixed' costs might increase in sharp increments if extra capacity is added.

Accounting data captures what is spent, where and when. To regulate consumption, we also need to attribute costs to cost objects, understand how they behave and how we can influence them.

Capturing and categorising cost data


CAPTURE



CATEGORIES



CHANGE



Cash, costs, and consumption – three different perspectives on resources.

But this isn't the full picture because 'costs' are not the same as cash (energy).

We can place the blame for this at the door of the profession responsible for recording and classifying flows of money in organisations – accountancy.

The bedrock of accountancy is the matching principle. This states that the value added (profit) is the difference between the resources consumed in providing goods or services (costs) and the income generated (revenue) for the period when the transactions took place. Accounting convention defines the reference period as a 12-month financial year (so called because it doesn't have to coincide with a calendar year). The financial year is usually subdivided into quarters or months (often 'book months' that do not coincide with irregular calendar months).

Because cash may not change hands in the same period in which the transaction took place, the profit for a period rarely agrees with the cash (i.e., resource) flow. The difference held in Balance Sheet accounts. The short term timing difference is called working capital, comprising debtors (what we are owed), creditors (what we owe) and inventory (what we have made but not yet sold). But if the transaction in question generates benefits over a number of years, the associated cost is treated as an asset that is charged to profit over its estimated life (depreciation or amortisation).

One of the challenges for accountants is establishing to what extent transactions impact the current year and placing a value on that part that is 'carried forward'. For example, the stock of unsold material can be valued in many different ways, all of which involve spreading fixed costs directly associated with the sale (e.g., manufacture or storage). Also, some tangible objects are treated as a (long term) fixed asset, but usually not that expenditure that helps create a long-term intangible asset, such as the advertising that builds the 'Coca Cola' brand. And we get ourselves in a real tangle with software. Is it an asset? For how long does it provide value? Does it matter whether we have purchased it or developed it ourselves? And is training people to use an asset part of the cost or not?

To be able to compute an accounting profit we cannot avoid the matching principle. But for the purposes of regulating resource consumption, we need to engage with the underlying reality, ignoring the impact of the effect of sometimes arbitrary accounting processes.

Effective regulation of resources requires that we look beyond the profit and loss account.

Cash, costs and consumption: three resource perspectives



Costing: attributing costs to 'things'

Matching costs with revenue is a start, but for decision making purposes we might need more information, which is where costing comes in. A costing process takes data captured in accounts and cost centres and assigns them to cost objects. A cost object is an entity of interest. These could be products (or services), customers, projects, or activities, that could then be aggregated into product groups, channels and so on. Cost objects can also be used to organise cost information so it can be matched with related revenue in what are called profit centres. There are three basic ways to attribute costs to a cost object.

The first is by direct **allocation**. Here the entire contents of an account/cost centre can be associated with a relatively 'large' cost object such as a project or a production site.

But if the cost of transactions relate to a range of cost objects, we need to spread, or **apportion**, the cost. This can work well (in the sense of being close to the 'truth') when the structure of a cost object is well understood, and we can use model-based apportionment. This is sometimes called standard costing.

For example, a manufactured product may be associated with a bill of material that specifies how much of what is used to make it, which can easily be costed by adding the estimated price of those inputs. Any difference between the predicted and actual costs (variance) can then be attributed to price or usage and apportioned pro rata to the 'standard' consumption. This is called marginal costing when the costs are fully variable, and absorption costing when the model is extended to cover fixed direct costs. In the latter case, we also need to estimate the volume over which the cost will be spread and any difference attributable to misestimated volume is called an over/under recovery variance.

Finally, simple apportionment is used where there is no direct causal relationship between a cost and a cost object. This uses an attribute of the cost object called an apportionment key (e.g., revenue, volume) to spread the costs.

In practice, costing processes often have multiple steps. For example, to avoid using arbitrary indirect apportionment keys like 'volume', accountants sometimes use a two-step process called Activity Based Costing. Here costs (e.g., 'salaries') are first allocated to an activity such as 'invoice processing' and then a 'cost driver' such as 'number of invoices' is used to apportion the activity cost to a cost object such as a product. This process is less arbitrary than simple apportionment and better reflects how costs behave, but at the cost of extra work and complexity.

Costs that can be allocated to cost objects or apportioned using a model are the best way estimate the true cost of 'things'. Indirect methods can be misleading.

Different ways of estimating the cost of 'things'



Cost ELEMENTS **Costing METHODS** **Cost OBJECTS**

e.g. materials


e.g. finished product


STRONG LINKAGE


ALLOCATE



DIRECT ALLOCATION



MODEL BASED



APPORTIONMENT

SPREADING



APPORTIONMENT KEY
(e.g. revenue)



ACTIVITY BASED

(secondary)



COST DRIVER
(e.g. # invoices)



WEAK(ER) LINKAGE

Stronger links
 more 'truthful'
 estimates of cost

Costing: Complexity quickly proliferates

To give some idea how tortuous costing can become let us work through a hypothetical example from a sample 'toy' company that has only four accounts (types of cost) and one cost centre. These costs need to be attributed to a set of 'product' cost objects (A1, A2, B1, B2) linked to two profit centres – A and B.

Initially the cost of materials (variable direct costs) is directly allocated to products based on a costed bill of material. Other production costs and distribution costs (fixed direct costs) are allocated to the profit centres A and B directly where possible and then apportioned (downwards) to individual products based on volume. Indirect costs (overheads) are apportioned to profit centres based on revenue but not spread to individual products. As a result, we can estimate the direct cost and gross profit for every product and a full cost and profit number for each profit centre.

Now suppose there are two distinct channels (X and Y) each comprised of two accounts, which are serviced in very different ways. To get a fix on customer profitability we need to calculate the 'cost to serve' more accurately. Fixed direct costs other than distribution can be picked up from the product costing. Distribution costs can be allocated to individual customers, but then need to be apportioned to products. And to attribute the costs of the sales function, invoicing and debt collection (currently lost in overheads) to customers we first must calculate the costs of these individual activities. Sales costs can be allocated directly to accounts, but invoicing costs are spread using the number of invoices and the average level of debt as cost drivers.

Even in this super simple example many different processes and apportionment keys are used, and some costs are allocated multiple times. The process is complex and result in two different versions of product costs depending on whether we are looking at the product or the channel dimensions. This procedure isn't necessarily wrong. We may need to allocate costs to manage profitability, and product costs can be used to set prices and customer specific costs to help set discounts. But the important question is, does such a 'sophisticated' costing process help us understand costs in a way that helps us to manage them?

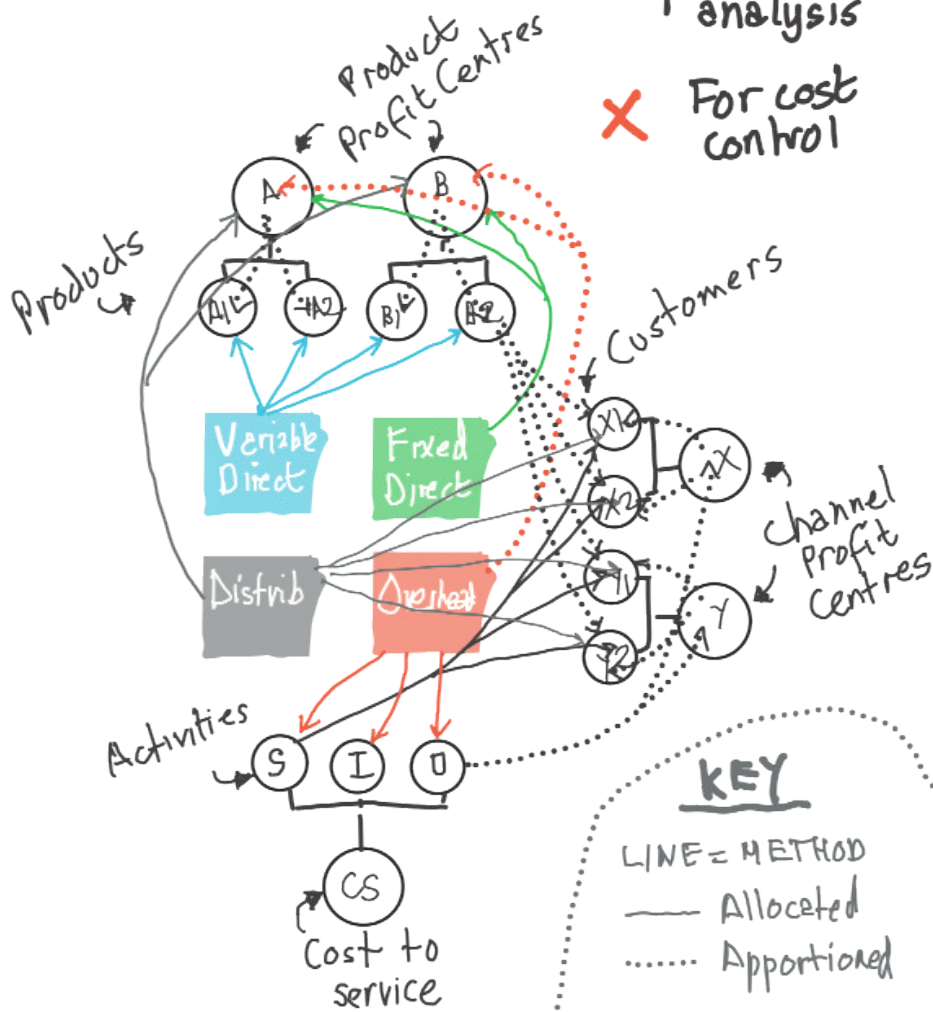
Other than in the case of truly variable costs, I think not. This is because it involves fabricating arithmetical links between costs and cost objects that don't exist in reality. It creates the illusion of causality and a false narrative about what is required to make improvements.

Costing helps manage profitability, but it is of limited value in managing costs.

Costing complexity proliferates

✓ For pricing & profitability analysis

✗ For cost control



ILLUSTRATIVE

* not all lines are shown *

The value of costing

From our review of costing processes, it is clear that we can never know the ‘true cost’ of a product or any other cost object. Any ‘result’ is the product of assumptions made by cost accountants to a greater or lesser extent. The cost of ‘x’ is an abstraction from reality, not the ‘truth’.

Every process should be designed to fulfil a purpose. I argue that the primary purpose of the costing processes I have described is to help manage profitability, not costs. But it is easy to lose sight of this when you are knee deep in the mathematical thickets of a process.

While we accountants make take pride in developing ‘sophisticated’ costing processes involving multiple steps and elaborate apportionment techniques, the reality is that they often obscure the causal relationships between costs and with the cost objects to which they are attributed. Indeed, they can make things worse because apportionment can create the illusions of a causal relationship between costs and cost objects that doesn’t exist in reality and so actively promote poor decision making.

Worse, I now recognise that I have often been complicit in using the costing process in an attempt to manipulate behaviour, for example, to ‘protect’ vulnerable parts of a business, or to force the management of another to improve their margins. And I have wasted weeks and months arguing about what is the ‘fairest’ way to allocated costs. Now I look back and I wonder what I was thinking. Clearly, we don’t want to do anything that grossly distorts the picture we paint of reality. But why should the perceived moral ‘fairness’ of the outcome be a criterion?

If our purpose is to produce information that helps us become more efficient and effective in how we consume resources, a costly, complex, and opaque costing process is more often part of the problem rather than the solution.

What other methods are used to manage costs? And how well do they work?

Purpose should inform process. Costing (rarely) serves the purpose of helping to regulate the consumption of resources.

Costing has a limited role...

... because

PURPOSE  PROCESS

Regulating
↓

Profitability  Costing

Resource Consumption  ?


So:

- ✓ Cost categories
- ✗ Costing

WHY?

Costing processes can be:

-  Complex
-  Costly

(A)  (B) Confusing
Cause Effect

Module 2: The existing structure

Preview

Having established how cost data is captured and processed we now turn to how it is typically used to control costs. There are three approaches traditionally used to regulate the flow of resources in an organisation. And while each follows a logic that fits well with traditional notions of control, they have weaknesses that can lead to resources being misdirected or wasted.

Responsibility accounting focusses attention on things that are easy to hold people accountable for (cost inputs) resulting in the systematic neglect of those things that they cannot (e.g. stocks). It can encourage the pursuit of false economies of scale and to excessive fragmentation in work systems. And it can lead to perverse outcomes when accounting conventions are applied inappropriately.

Budgeting is a process that aims – amongst other things – to control costs by imposing a system of fixed constraints on spend. But it is administratively burdensome to administer, which makes it part of the cost problem. The constraints set often restrict an organisation's ability to respond to unexpected threats or opportunities. They can also generate misleading information and promote dysfunctional behaviour that locks in waste and leads to perverse decision making.

In addition, **authorisation processes** are frequently used where the level or nature of a resource allocation decision is deemed to require extra control. But the way that this is often applied falsely assumes that more senior people always have the knowledge and expertise required to make good decisions, which may not be the case.

Also, in conditions of great uncertainty (e.g. where a long term commitment is required or a major innovation), forcing things through a process where a binary yes/no decision is required before resources are released is problematic because by so doing it removes the ability to change in response to new information. In other words, it closes down options. And, as I discovered when this was a major part of my job, in uncertain conditions it is easy to produce credible arguments for any number of potential outcomes, but difficult to establish which one is 'best'. As a result there isn't one number upon which you can safely base a decision.

All these problems arise because the regulatory systems have been designed around structures constructed for other purposes, such as organisation charts. And they are based on the assumption that workers in the system cannot be trusted, so control has to be imposed on them from outside.

The solution is to build the regulatory process into the work system and to give people within it the ability and the tools to control it for themselves.

The traditional three-pronged approach to cost control

The traditional approach to controlling costs is based on three pillars.

The first is called RESPONSIBILITY ACCOUNTING. The principle behind this approach is that every element of cost (and for that matter revenue) should be controlled by one (named) individual. If they are successful in their efforts, they can be rewarded, if they don't, they can be 'held to account' (i.e., punished in some way).

Accountability can be based on accounts (cost type), cost centres (groups of accounts), cost objects (an entity to which costs are assigned) or a combination of all three. And because roles are usually arranged hierarchically, accountability can be exercised at multiple organisational levels. So, a boss can delegate some or all of his or her responsibilities, for example.

The second pillar is BUDGETING. Whereas responsibility accounting establishes the structure of accountability, budgeting determines what an individual is held accountable for. In the case of costs, a budget limits spending to a predetermined fixed amount for a given year (or quarter or month), cost centre (and/or account) or cost object (or cost within a profit centre).

The third pillar is AUTHORISATION. In addition to allocating responsibility for a set of costs and setting a limit on what can be spent, some purchases or projects may also need to be authorised or 'signed off' by someone – usually someone more senior to the person spending the money.

What type of expenditure needs additional authority and from whom often depends upon the size or nature of the commitment (e.g., any new hires, all purchase orders or projects in excess of £n,000). And because the cost of assets (depreciation) is spread across many years (thereby diluting the impact of control based on the profit and loss account) capital expenditure will normally be subject to special authorisation procedures.

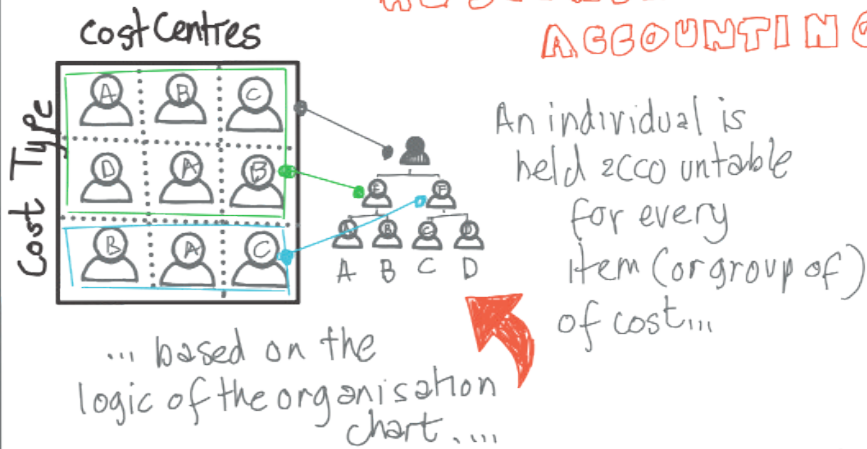
In summary, the traditional approach to controlling costs depends on making someone accountable and then imposing constraints on what can be spent and when.

The question is: how well does this work?

Traditional cost control is based on the principles of responsibility, budgets, and authorisation.

Three pillars of traditional cost control

RESPONSIBILITY ACCOUNTING




ANNUAL BUDGETING

... which in total can be no higher than the annual budget...



AUTHORISATION

... and may require additional authorisation from "someone higher" in the chain of command.



Responsibility Accounting: Sub optimization trap #1 – the invisibility of stocks and waits

The key argument for responsibility accounting is that every item of cost can be assigned to a single individual who, if suitably incentivised, will constrain the cost of the parts that they ‘own’ and so reduce the cost of the whole system.

The flaw with this argument is that resources are usually consumed by a work system involving many parts of the organisation. By subdividing it for the purposes of control we risk losing sight of the relationships between the parts that generate costs thereby suboptimizing the performance of the system, particularly when we use the structure of the organisation chart, not that of the work system itself.

Let me illustrate this with an example using a simple ‘toy’ process comprised of three steps – ‘owned’ by individuals A, B and C reporting to bosses X and Y. Let’s assume that each ‘owns’ a cost centre for which they each are held responsible – as measured by total cost or cost per unit. In between each of these steps there is a gap. In the case of a manufacturing process the gap could be made up of stock or work in progress – the result of the individual processes working at different rates in an uncoordinated way. If this were a service process, we might see a queue (a wait).

While each of the steps, and the costs associated with them, is clearly assigned who is responsible for the resources tied up in between each of the steps? The answer is usually no-one, because of the way the system is set up. And this situation is exacerbated by the way that physical stock is accounted for.

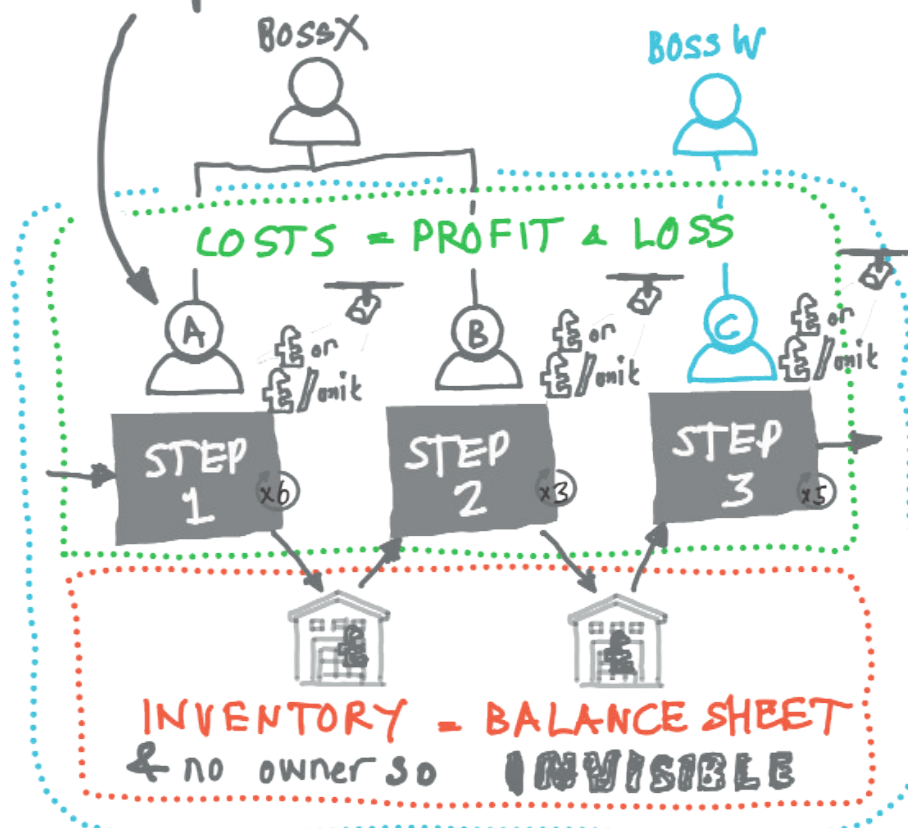
The resources consumed by the three processes are treated as costs in the profit and loss and the rate at which they are consumed can be easily measured. And it is clear that less (cost) is better than more. On the other hand, stock is held in the balance sheet and so isn’t treated as ‘current’ consumption. Furthermore, it isn’t always clear what a ‘good level’ of stock is. Too much stock is bad because it is unproductive, but if there is too little the receiving process might run out of work to do (or customers might not have their needs met) which is clearly bad.

And the cost of too little or too much stock is difficult to pin down. For example, the cost of poor service or lost sales doesn’t appear anywhere in the accounts. And the indirect cost of holding too much stock – deterioration, damage and wasted space – often cannot be easily traced and assigned to a single cost centre. So, in practice, no-one is held accountable for a key aspect of system performance and problems are either rendered invisible or kicked down the road.

Optimising the parts of a work system sub optimises the whole – often reflected in levels of inter process stocks or queues that no-one is responsible for.

Suboptimisation trap #1 the invisibility of inventory

By making individuals responsible for the costs of their parts...



... you suboptimise the whole system