

MAKING ENDS MEET: A CONCEPTUAL AND ETHICAL ANALYSIS OF EFFICIENCY

ABSTRACT

Efficiency is often overlooked as an ethical value and seen as ethically relevant chiefly when it conflicts with other values, such as equality. This paper argues that efficiency is a rich and philosophically interesting concept deserving of independent normative examination. Drawing on a detailed healthcare case study, we argue that making assessments of efficiency involves value-laden, deliberative judgements about how to characterise the functioning of human systems. Personal and emotional resources and ends are crucial to system functioning, but often discounted in favour of a relatively narrow set of financial inputs and institutional or procedural outputs. Judgements about efficiency tend to advantage (or disadvantage) different parties, depending on the resources and ends considered. Different constructions of efficiency can therefore promote or neglect the perspectives and interests of differently placed actors. Models of efficiency do not merely embody contestable ethical standpoints but—put to use—can unwittingly reify and reproduce them.

KEYWORDS

efficiency; value; deliberation; measurement; healthcare administration; management

1. INTRODUCTION

In comparison to obviously normative and contested values, efficiency might appear to be relatively straightforward. Specifying what it means for an institution or action to be fair, or legitimate, or equitable is liable to generate immediate disagreement, not only about the relative importance of these ethical values but about their meaning. But defining and identifying efficiency may, at first sight, seem less controversial: when we have the facts and figures in front of us, it is difficult to dispute whether one course of action is cheaper, or quicker, or otherwise uses fewer resources than another to achieve a similar end. Efficiency is most often treated as an ethical concern when it threatens to come into conflict with other values (Culyer 2015). In this paper, we turn an analytical lens on efficiency, seeking to unpack some of the conceptual and ethical complexity that is involved in generating models and making assessments of efficiency that are nuanced and fit-for-purpose. We seek to show that efficiency is a much more philosophically rich concept than is typically realised and that it deserves a detailed normative analysis in its own right.

It is not necessary to scratch far beneath the surface to appreciate that efficient systems are not merely those that choose the quickest and cheapest options. In the public sector, austerity politics and funding cuts are often justified in the name of efficiency, with their advocates leaning on metaphors of waste reduction and ‘trimming the fat’ in order to rationalise shrinking budgets and disinvestment in public services. But these can lead to substantial failures in the delivery of high quality, sustainable service and sometimes catastrophic longer-term outcomes for the people who use public services and the staff who work in them. In the UK, for example, lack of investment in health and social care, and particularly in public health

and health security, as part of sustained austerity drives arguably contributed to a weakened health system that lacked the resilience required to respond adequately to the Covid-19 pandemic (Merry and Gainsbury 2023). One way of understanding such problems is that they result from efficiency being unduly prioritised above other important social goods and values. But, in a keynote speech given in the UK Parliament, the head of the National Audit Office argued that public sector efficiency requires *more* investment in health services, the courts, the Passport Office, and government digital services, not less, in order to improve operational efficiency (Davies 2022). This framing suggests a different way of understanding the issues associated with over-zealous efficiency campaigns—the problem is not over-emphasis on efficiency but a mistaken understanding of efficiency. Efficiency isn't just a matter of cutting funding and resources, but rather requires judicious and well-justified use of resources, in order to best promote and secure the things that matter. In this paper, we set out a more extended case for this second picture of efficiency: a value that is concerned with balancing and optimising things that matter, rather than penny pinching and parsimony. Understood thus, determining what is efficient requires careful reflection on the range of things that matter and how to best pursue and prioritise them, as well as broader consideration of the goals and function of systems and services.

We begin by showing that the concept of efficiency can be construed and applied in both more narrow, technical ways, and in broader, evaluative ways. Narrower constructions are concerned with merely minimising waste; broader constructions are better understood as being about *optimising* inputs in relation to outputs. In relation to human systems and social institutions, understanding what efficiency means for a given system requires the relevant set of resources and ends to be identified and defined. In practice, making efficiency assessments in relation to social organisations and human systems requires deliberative judgement about a dynamic and interacting complex of factors.

We go on to show that efficiency is not just conceptually complex, but also normatively complex. Assessments of efficiency are perspectival; they are made by someone (or some institution), about some system, for some purpose. Moreover, different ways of framing efficiency, and the resources and ends that characterise it, will benefit and disadvantage particular actors and institutions differently. Decisions to define efficiency in a particular way are not just decisions about what resources and what ends are seen as relevant, but also decisions about *whose* resources and *whose* ends are treated as relevant. Different constructions of efficiency can therefore promote or overlook the perspectives and interests of differently placed actors. We argue that there is good reason both to be open, in theory at least, about what the relevant resources and ends are, and to recognise whose resources and ends they are, even if in practice it is difficult or even impossible to include them all in a useful model of efficiency. In resource-constrained systems, it is inevitable that difficult decisions will have to be taken about how to define efficiency and about what ends to prioritise at any given point in time. But we suggest that identifying and acknowledging the wider implications of such decisions can ensure that efficiency claims are nuanced and do not overlook the full costs, including human costs, of low-(financial)-cost healthcare.

Finally, we consider some of the ethical implications of modelling and assessing efficiency in practice. Models of efficiency are not just passive descriptions: the subsequent

use of such models to manage resources has the potential to shape services in ethically significant ways. In practice, actions or strategies designed to improve efficiency are liable to be quite technical in nature, focusing on relatively circumscribed resources and ends. Such improvement efforts, however, will only ever reflect a partial, value-laden construction of efficiency. We therefore suggest that efforts to improve technical efficiency should incorporate deliberative reflection on broader constructions of the resources and ends that characterise services—even when it is not possible to measure and operationalise these resources and ends directly. Assessments of technical efficiency should be seen as a tool to be used in deliberative decision-making about efficiency. While they may appear to provide clear answers about what is most efficient, they offer limited and partial accounts, and their apparent authority can obscure the ethical choices that underpin them.

The complex nature of efficiency is made manifest when looking at the detail of real-world cases. Our discussion draws on a case study of a London-based operational research project—WORKTECC (Workforce Operations that Realise Knowledge-based Transformational Efficiency gains in Community Care)—which aimed to characterise and model efficiency in home healthcare services. This case offers many insights about efficiency in relation to home healthcare, but also illuminates the concept of efficiency more broadly, helping to expose the value-laden decisions that underpin claims about efficiency—in the healthcare context and beyond. As well as providing a more general conceptual and ethical argument about efficiency, the present paper contributes to the regrettably underpopulated literatures on the ethics of healthcare management and public administration.

2. MAXIMISING AND OPTIMISING

We begin by introducing the concept of efficiency. We go on to distinguish between *technical* efficiency, which is concerned with maximisation, and *productive* and *allocative* efficiency, which are concerned with optimisation.

Efficiency, very broadly speaking, is an assessment of a system that captures its ability to maximally achieve its ends with a minimum expenditure of resources (Alexander 2009). It is largely silent about the substantive content of the ends and resources in question, and instead concerned with how to best manage resources in pursuit of those ends. Efficiency is, thus, in some sense ‘empty,’ needing to be filled with other values—without specified ends in view, the most efficient thing to do would be to use no resources and to do nothing. But efficiency is not thereby value-free and is closely related to other procedural values such as maximisation and optimisation. This suggests that the value of efficiency is related to the opportunity costs of inefficient functioning in conditions of scarcity. That is, a system that is not maximally efficient *could* produce more than it does, given the resources it operates with and the constraints that it faces—or could produce the same output with fewer resources—and, other things being equal, it would be better were it to do so. Efficiency is used to characterise the functioning of mechanical systems—typically by calculating the ratio of power output to power input—but it is also used to characterise the functioning of economic and social systems. In these contexts, assessments of efficiency can potentially draw on a wide range of resources as inputs—money, time, personnel, equipment, energy, and so on—and consider a wide range of

ends as outputs—such as provision of products, services, knowledge, qualifications, public goods, and production of personal utility, wealth, and health.

At first glance it is quite easy to understand why efficiency is valued in human endeavours. If you can achieve the same desirable outcome using fewer resources or achieve more of what you want using the same resources, so much the better. In business, greater efficiency typically equates to greater dividends to shareholders or more investment back into the company. In the public and non-profit sectors, greater efficiency typically equates to more social goods or fewer opportunity costs, that is, more resources to be used on other important activities. But efficiency is not necessarily equally valuable in all fields. It might not make much sense or be desirable to cultivate efficient friendships and relationships or increase the efficiency of conversations. And there may be limits to the sense in which it is valuable for art or literature to be efficient. Nonetheless, in relation to more instrumental endeavours—such as delivering public services, providing and distributing humanitarian aid, and running a business—efficiency helps people to do more of what they want. Despite this pro tanto value, it is important to note that efficiency can be used for ill as well as for good. For example, some have argued that the trans-Atlantic slave trade and management of American plantations were efficient compared to farms worked by free labourers, and that plantation owners were thus rational, maximizing agents (Fogel and Engerman 1995). While the actual efficiency of slavery as an economic model is disputed (Gutman 1975, Hilt 2020), a widespread belief in its efficiency could plausibly be used to support arguments in its favour. Another example of misdirected efficiency is nineteenth century engineer Fredrick Winslow Taylor’s system of scientific management, which broke down factory labour processes into precise, repetitive motions in an effort to increase the output of workers, and in so doing treated human labourers as machines, exploiting and mistreating them physically and psychologically so as to increase the profit-margins of their employers (Scott 1998, Alexander 2008). These examples highlight that the ethical value of efficiency cannot be disentangled from the ethical value of the aims of the activity in question: if the aims are morally compromised, it won’t make it more ethical to achieve them efficiently. Indeed, it might be more *unethical* if and when efficiency is used to promote and advance unethical ends.

Efficiency is widely invoked as a core healthcare value: a good healthcare system is one that does not waste its resources in pursuit of its ends. In principle, the resources in question are defined very broadly; the US Institute of Medicine, for instance, suggests that an efficient system avoids wasting ‘equipment, supplies, ideas, and energy’ (Institute of Medicine (US) Committee on Quality of Health Care in America 2001). In practice, however, efficiency is often understood chiefly in relation to a relatively limited set of resources, such as financial cost, time, equipment, and workforce (Hussey, De Vries et al. 2009, Varabyova and Müller 2016, Cantor and Poh 2017, Zakowska and Godycki-Cwirko 2019). And while there are many possible ways of characterising the ends in question—including both outcome and process measures—typically efficiency at the level of healthcare institutions is understood in relation to service related activities rather than health outcomes (McGlynn and Shekelle 2008, Hussey,

De Vries et al. 2009, Cantor and Poh 2017, Zakowska and Godycki-Cwirko 2019).¹ Efficiency is also invoked in many other social and institutional contexts: the UK Government Efficiency Framework seeks to guide “every public servant involved in delivering our public services” in making sure they are maximising output for the minimum input (Government Finance Function 2023). Here the focus is firmly on spending money wisely, and other resources are considered chiefly in relation to their financial impact. In the UK higher education sector, *Efficiency and effectiveness in higher education*, a report by the Universities UK Efficiency and Modernisation Task Group almost exclusively focuses on financial cost, and in particular the need for universities to reduce costs while improving the quality of their teaching, research, and front-line services (Universities UK Efficiency and Modernisation Task Group 2011).

Good social systems and institutions are rarely *merely* efficient. Efficiency is typically considered alongside, and constrained by, other values—such as effectiveness, fairness, responsiveness, sustainability (Maxwell 1984, Institute of Medicine (US) Committee on Quality of Health Care in America 2001, Mortimer, Isherwood et al. 2018). It is tempting to see these other values as inherently challenging to conceptualise and measure, while efficiency is their more straightforward cousin—simply requiring that these other values be achieved at minimum cost. We suggest that this temptation should be resisted and that efficiency is not so straightforward. We also suggest that efficiency is importantly unlike other values, insofar as it plays a mediating role between them. That is, a social system is efficient to the extent that it optimally balances and reconciles its (potentially wide-ranging) commitments and ends. Making a system more efficient is itself an activity that uses resources, so efficiency can also be treated as one goal among others—and, indeed, may need to be balanced against other ends (Sassi, Le Grand et al. 2001). We illustrate with examples relating to healthcare, but endeavour to draw more general conclusions about the value of efficiency.

It is possible to characterise efficiency in a *technical* way (Cylus, Papanicolas et al. 2016). Given a system with a clearly defined and stable goal and a clearly defined set of resources, relatively uncontentious claims can be made about its efficiency. A technically efficient system will achieve the desired outputs using minimum resources or maximise the desired outputs using a given input of resources. So—using an example from the healthcare context—assuming a fixed list of patients who need to be visited at home by a community nurse on a given day, it will be more-or-less straightforward to minimise travelling time between home visits. Depending on whether the input (time) or output (number of patient visits) is treated as fixed, an efficient system might either maximise the number of patient visits within a defined period or minimise the number of hours used to visit a defined number of patients.

But, in practice, even relatively simple descriptions of services such as home healthcare include multiple inputs and outputs. Such services operate with diverse resources—not only time, but also, for example, a relatively fixed number of staff, with a range of professional competencies, and a stock of equipment. They also operate with a series of ends—not just

¹ Population level health outcomes such as life expectancy are used as measures of output when the efficiency of entire national or regional health systems is under consideration (Varabyova and Müller 2016).

visiting patients, for instance, but meeting their clinical needs and satisfying patient preferences in relation to timing of visits and continuity of care. As soon as a number of variables are in play, it becomes difficult to make unambiguous claims about efficiency (Alexander 2009). Scheduling the maximum number of visits might involve not meeting some patients' preferences about the timing of visits; assigning staff to patients in order to maximally meet their clinical needs might go against patient preferences about continuity of care. The variety and interaction of factors that characterise a home healthcare service, as well as any other social organisation or human system, even on a cursory reflection, suggests that it is not sufficient to conceive of service efficiency as a scalar quantity obtained from a technical input-output calculation.

The concepts of *productive* and *allocative* efficiency help to capture the way that efficiency in complex systems represents a balancing of different ends, constraints, resources and commitments. Productive efficiency captures the minimum costs needed to achieve a given output, or the maximum output that can be achieved for a given cost (Palmer and Torgerson 1999, McGlynn and Shekelle 2008). Calculating productive efficiency involves converting the different resources a system uses into a common value—typically financial cost—in order to calculate the overall input cost. Productive efficiency does not require the minimisation of all inputs, but rather represents the optimal mix of inputs for cost minimisation. Productive efficiency enables simultaneous consideration of a variety of different inputs and so captures more than technical efficiency. In practice it also requires all of the inputs to be represented in financial terms. This makes it rather a blunt instrument for thinking about efficiency, particularly in relation to human systems which can be characterised in terms of resources that are not easily attributed a financial value.

Allocative efficiency captures the optimal distribution of outputs for a system (Williams 1988, McGlynn and Shekelle 2008, Cylus, Papanicolas et al. 2016). A system is allocatively efficient if, given some set of inputs, it is not possible to increase or improve any output without decreasing or worsening another (Le Grand 1990). Allocative efficiency therefore does not necessarily maximise any individual outputs—for maximising one output might relatively diminish another—but rather *optimises* across a series of outputs. In economics allocative efficiency is often defined in terms of individual utility—where allocative efficiency indicates that making any individual(s) better off will involve making some other individual(s) worse off (Knight 1992, Palmer and Torgerson 1999). When applied to home healthcare services, the inputs and outputs in question are perhaps more likely to be understood in institutional, rather than individual, terms—that is, allocative efficiency in this context involves optimising institutional ends given a range of resources.² This is, in part, because it will be difficult to convincingly convert all the relevant outputs of a healthcare service into the currency of individual utility gains.

² We will, however, go on to suggest that a nuanced understanding of institutional resources and ends, and ultimately of efficiency, requires consideration of the relationships between institutions and individual staff, patients or service users, and other stakeholders.

Technical, productive, and allocative efficiency are all ways of characterising efficiency. The move from *maximisation*, which is reflected in technical efficiency, to *optimisation*, which is captured in productive and allocative efficiency, reflects recognition of the fact that in even a moderately complex system, when it is difficult to adequately capture resources and ends using a single value currency, maximising the achievement of those ends and minimising the use of resources is a balancing act. There is unlikely to be a course of action that simultaneously improves the system with respect to the achievement of all desired ends and the reduction of all resources (Van de Poel 2009). Improvement in some respects might, then, involve compromise in other respects. Allocative efficiency in particular is a helpful tool for characterising and assessing efficiency in complex social systems because it captures both the irreducible variety of ends served by such systems and also reflects the importance of going beyond simply thinking about cost reduction in serving those ends. Allocative efficiency represents a more evaluative conception of efficiency; there is frequently not a clear single most efficient option, but rather several, equally efficient options that differently satisfy the array of ends using an array of different resources.

Efficiency in complex systems like healthcare institutions, requires something more like a judgement rather than a calculation. Two further factors complicate the assessment of efficiency in such systems: i) the inputs and outputs of a system will be in some respects relatively fixed but also somewhat flexible and shifting in extent and nature over time and ii) the relevant set of inputs and outputs, that should be used to determine efficient functioning, is to some degree open to debate. We illustrate each of these factors with examples from the home healthcare context.

First, there will be some degree of flexibility in the definition of the ends and resources of a home healthcare service. The availability and use of some resources may be—at least partially—subject to the service’s discretion, while other resources will be more fixed. Some of the desired outputs of a service will be more critical, others more dispensable. So, for instance, contractual constraints might limit the income that a service operates with and the services that it is required to provide, though there will be some leeway in how the money is spent and services are delivered. Legal constraints may more rigidly determine the rights and obligations of its workforce and patient population. And while such a service will, to some extent, be able to shape its human resources through recruitment of staff and through training, its ability to do so will be limited by its income stream and the broader social context of healthcare education and the labour market. Similarly, while a service may be able to shape its outputs to some extent, for example by delaying patients’ entry into the service during busy periods, its scope for action will be heavily influenced by external factors.

Moreover, many of the required ‘outputs’ of a system will be in flux. Patients have diverse needs, have different levels of existing informal support, require different kinds and degrees of medical and social support, and all these may change over time. This will affect the necessary frequency and length of visits. There are constraints on when home visits can occur, for example, if patients attend day centres or other scheduled social arrangements, or take holidays. Over time, the set of patients changes, as new patients are referred into the service and others discharged. The set of inputs will also be in flux. The total number of working hours available for home visits depends on the working hours of staff, which may change from week

to week, conditional on training, sickness, holidays, administrative work, and other commitments. Aggregate working hours are also affected in the longer term by factors such as retirement, recruitment, maternity leave, and changing work patterns. The staffing skill mix will change with staff turnover, training, promotions and accumulation of experience over time. Inputs such as equipment and supplies, and of course the funding allotted to the service, may be influenced by broader economic factors and by more proximate decisions about service design. While allocative efficiency retains the core idea that efficiency involves doing as much as possible with the available resources, it must involve consideration not just of a fixed set of inputs and outputs, but of the space of possibility of those inputs and outputs—how they can change and the knock-on effects of such changes.

Second, in addition to the most commonly invoked set of inputs and outputs—time, money, staff numbers and skill mix, service user numbers and needs, visits and interventions performed—there are other, more complex, human factors that are relevant to thinking about organisational efficiency over the longer term. In the home healthcare context, clinical outcomes—such as admission to hospital, worsening symptoms, complications, and death—and operational outcomes—such as waiting times, total number and average length of visits—might be the most evident ways of assessing whether a service has achieved its ends. But the ‘outputs’ of such a service will extend beyond this. Patients often have other concerns in addition to those relating to clinically appropriate management of their health, including preferences about continuity of care, clear communication, and the social and interpersonal aspects of care and conversation. The achievement of such outcomes is a matter of degree and there may not be a clear threshold to designate success, nor means of measuring it. Factors such as staff retention and staff well-being could, at least in principle, also be considered goals of a home healthcare service in their own right, as well as means to other ends. In any public service, the ‘inputs’ will also extend beyond time, money and staff numbers, as the ability, productivity, and commitment of staff are linked to factors such as workload, team composition, workplace support and management, appropriate use of their skills, and their well-being and motivation more generally. Assessing staff productivity and well-being arguably includes taking seriously the agency, virtues and particular contributions of individual staff members. Failure to attend to these factors could lead to high staff turnover in the longer term, as well as additional costs associated with recruitment and induction, the employment of temporary workers, staffing gaps, and so on. But, moreover, such a failure may also impact on system outcomes as well as on resources—affecting, for instance, the quality of services and communication experienced by users, outcomes, and the number and gravity of errors made. If staff believe they are treated as fungible resources as a result of efficiency drives, this may create a vicious circle.

These complicating factors together suggest that making assessments of efficiency requires substantial deliberation about a range of factors, and may require holistic, systemic modelling of the complex ways in which resources and ends intersect and influence one another, as much as granular, linear modelling. They also suggest that a reasonably comprehensive understanding of efficiency in such contexts requires an understanding of institutional and personal variables and the ways these interact. To overlook the ways that institutional decisions about resources impact on and are impacted on by the personal resources and ends of individuals is to operate with a highly limited conception of efficiency.

Adequately modelling the causal relationships between personal and institutional variables is fiendishly difficult, however. For example, predicting how changes to workload allocation will affect staff retention or short- or long-term outcomes would require taking into account the huge range of other inter-relating variables that are impacted by, and impact on, these factors. Modelling this could draw on behavioural psychology and behavioural economics, for example, to estimate the impact of tiredness and stress on service quality, decision-making and communication, as well as more standard considerations in operational research such as time, workforce capacity and skill-mix. But in many instances, not least in relation to the effects of care provision on long-term health outcomes, there is simply no credible knowledge base to draw upon. Using clinical and personal outcomes, such as individual well-being, as outputs for assessing efficiency may appropriately reflect the generally recognised purposes of a service, but involves methodological challenges: it depends on non-arbitrary ways of defining the time-frame over which outcomes are measured and accounting for the impact of non-healthcare-related factors on those outcomes (McGlynn and Shekelle 2008). In practice assessments of efficiency and proposed interventions to improve efficiency are liable to have a more technical character, operating with a smaller range of resources and ends, which are easier and less controversial to identify and measure.

We will say more about the relationship between more technical assessments of efficiency and efficiency more broadly understood in the final section of the paper. But, for now, it is worth noting that technical efficiency, whilst easier to model, is not insulated from the complexities explored in this section. Rather, calculating technical efficiency involves the construction of simplified scenarios, focussed on particular inputs and outputs, and thus setting aside other factors that might be salient to understanding efficiency. Specific measures of technical efficiency operate within a broader evaluative notion of efficiency, which makes it possible—indeed necessary—to question whether interventions that are judged to improve efficiency using such measures *do in fact* improve efficiency.

Defining and measuring efficiency in a social system is a complex matter. It is, moreover, a *normative* matter. Optimising the use of resources involves deliberation about what the system is and does, including consideration of its ends and what is needed to achieve them (Van de Poel 2009). Different emphases in the description and relative valuation of ends will open the door to different ways of optimally deploying resources in service of them. There is no incontestable fact of the matter about what any complex, human system is *for*. A system like a home healthcare service plausibly has several purposes and at least one of its central purposes—relating to the promotion of health and wellbeing—can be unpacked in several different ways. Different framings of efficiency emphasise different goods and ends and balance them differently against one another. In the next section, we will focus on one sense in which characterisations of efficiency have significant ethical consequences, namely, insofar as they reflect, or fail to reflect, different perspectives in their characterisation of system functioning.

3. WHOSE RESOURCES? WHOSE ENDS?

To illustrate the ethical issues arising from attempts to assess allocative efficiency, we will draw on insights from the WORKTECC project, which engaged extensively with

stakeholders in order to characterise and model efficiency in two home healthcare services in North-East London (Grieco, Utley et al. 2021). The project sought to identify commonalities and tensions between different perspectives on efficiency and, drawing on this, to model potential efficiency gains while explicitly acknowledging the trade-offs these involve. While this was chiefly an exercise in mathematical modelling, the stakeholder engagement and sensitive and nuanced conceptualisation of efficiency underpinning the production of the models generated a rich seam of qualitative evidence, which exposes the normative, as well as conceptual, complexity of the task.

Engagement with a variety of stakeholders revealed a wide array of characterisations of the actual function of the home healthcare services in question—and, furthermore, what good and efficient functioning could look like. Inefficiencies identified by staff included unnecessary travelling time, duplication of visits or tasks, the use of highly-skilled staff to perform tasks that could be performed equally well by lower-skilled—and lower-paid—staff, and delays in referral and assessments leading to clinical decompensation. From an institutional perspective these may look like wasteful use of resources, which could, and should, be remedied. In many cases, however, what appear to be inefficiencies from one perspective can be characterised as efficient functioning from another. For instance, taking patients' commitments into account with respect to the timing of their visits might result in additional travelling time for nurses. Similarly, satisfying patients' preferences about continuity of care—which providers sometimes also see as highly salient, for example, in home mental healthcare—could require highly-skilled nurses to sometimes perform low-skilled work, if, for instance, a patient's clinical needs change over time. Optimising merely with respect to travelling time or professional competency would prioritise the service perspective over that of the patient population.

These findings indicate that it is not just that optimisation involves prioritising certain valued ends and resources over others, but that sometimes the trade-offs entailed implicitly prioritise the resources, ends, and perspectives of some stakeholders over others. Of course, in practice, the picture will be far more complicated than we have suggested here and the trade-offs in question can highlight tensions within, as well as between, different perspectives. Continuity of care may, for example, have a series of communicative and ultimately clinical benefits (Pereira Gray, Sidaway-Lee et al. 2018). And optimising with respect to professional competency might have economic costs for the service—if specialised staff perform only the tasks for which they alone are qualified, further visits by less qualified staff may be required to meet a patient's clinical needs.

Stakeholder engagement also indicated that the inputs and outputs relevant to characterising efficiency are far broader than are typically taken into account. Patient preferences have acquired a relatively secure position for assessing the success of healthcare (Santana, Ahmed et al. 2019), so it is unsurprising to see them playing a role in efficiency considerations. But the WORKTECC project highlighted a range of personal resources and ends that are more rarely emphasised. Deploying a specialist nurse to perform tasks that could be carried out by a healthcare assistant can represent a financial inefficiency and also a clinical inefficiency, if it results in skilled care being delayed or unaddressed elsewhere in the service. But failure to optimise in this respect may also have personal costs for the staff in question. A

skilled nurse who is consistently required to perform tasks that do not make use of his skills may suffer personal and emotional costs. He may feel demotivated or unappreciated and while a certain amount of ‘good will’ may enable him to work below his skill level some of the time, or for a period of time, this is likely to be finite. Mismanagement of these personal resources and ends—motivation, good will, emotional energy, well-being, self-worth, dignity—can, of course, be understood as inefficient in institutional terms. Lack of practice at specialist tasks, depleted motivation, and low well-being may impact on a nurse’s ability to do his job well, contributing to worse patient outcomes or administrative errors. It may also make him more inclined to leave his job, which would carry additional recruitment and training costs for the service. But even if such personal misfortunes are *not* instrumentally detrimental—that is, clinical and operational outcomes do not suffer as a result—there is still reason to see them as costs in their own right. Human resources comprise not just staff numbers and skill mix, but also the personal and emotional resources available to staff, which play a role—potentially a very important role—in enabling human systems like home healthcare services to achieve their ends.

This case study highlights the way that efficiency gains and losses can accrue to different parts of a system and to particular stakeholders. To illustrate further, consider alternative ways of making use of efficiency gains. If more efficient route planning and scheduling of home visits were to lead to substantial savings of staff time, a further question arises as to how the ‘spare’ time is used. It might, for instance, enable longer visits, or more visits per day. Or it might be used to give nurses time to complete their administrative duties during their contracted hours, relieving them from having to work overtime or outside of working hours to do their job to their satisfaction. Or, instead, it might be used to give nurses a lunch break. Using the efficiency gain in any one of these ways will have efficiency implications elsewhere in the system, as it has opportunity costs with respect to other resources or ends. Such decisions and trade-offs are normative in nature. This is not just in the sense that any public services resource-allocation decisions are normative, insofar as they promote some ends at the expense of others, but also in the sense that they concern the ways that the interests and ends of different people interrelate and conflict. These are complex, ethical decisions, which reflect consideration—or lack thereof—of the personal ends and interests of different stakeholders. Such decisions have the potential to promote—or fail to promote—the personal agency of individuals, both in relation to institutional ends and interests and in relation to one another.

So far, we have considered how different constructions of efficiency might impact on different individuals in a personal capacity. But constructions of efficiency are perspectival in multiple senses: as well as being made by some individual or institution, they relate to some service or entity (identified and characterised in a particular way), and they are made with some objective in mind (McGlynn and Shekelle 2008). Different ways of framing efficiency privilege different institutional perspectives and social ends. In the healthcare context, assessments of efficiency at a broader system level can look different from assessments made at a more localised service level. One of the functions of home healthcare, for example, relates to keeping people out of hospitals—whether that be through enabling earlier discharge, preventing hospital admissions, or avoiding readmissions. This function can be conceived in

terms of system efficiency: the costs of home healthcare can represent, all things considered, savings on hospital admissions. This system-level efficiency focus might justify putting or transferring more resources into home healthcare services. But the fact that something is efficient for a system does not override priorities as seen from within a service perspective. Within the service, a system level goal such as reducing hospitalisation will comprise just one of a number of parallel outcomes that must be balanced and managed and sometimes it will be relatively deprioritised.

What counts as efficient functioning may look different from different vantage points within a system or service. Efficiency savings within a service could be used to expand its size or scope—that is, enable it to do more with the same resources—or instead, the freed-up resources could be moved elsewhere in the system. From the service perspective, neither of these options straightforwardly represents a more efficient course—either would optimise resources with respect to ends, but they would affect different ends and resources. From a system perspective, however, one option may be more efficient, for example if an expanded or a different service better supports the ultimate purposes of the system. So, for example, given a broader system that was aiming to move away from a bio-medical model, where hospital-based clinical practice is seen as the archetype of healthcare delivery, there could be a crucial place for an expanded and extended home healthcare system. In such a system, health would be understood less in terms of clinical outcomes and more through a social model, which views health as part of well-being and healthcare as something to which many different actors, in a wide variety of settings, can contribute. Within this imagined context it might be more efficient for a system prioritising these more holistic ends to expand and reorient its home healthcare services, rather than to move resources to secondary care or elsewhere. But even given an ideal of efficient functioning at a system level, achieving an optimal balance between resources and ends in its services remains difficult. In the home healthcare context, for example, balancing resources between responsive ‘crisis teams’ that step in when patients have acute needs and the longer-term community teams that provide more sustained and routine care is challenging. Shifting resources from crisis care to more stable, long-term community care may be effectively impossible given the ongoing urgent needs of patients, even if this would represent a more optimal use of resources from a system perspective and when thinking about a broader, longer-term view of service aims. What is best characterised as ‘efficient’ will reflect a vision of what a given system is supposed to be doing and characterisation of its ultimate goals.

Each of the ways in which efficiency is perspectival highlights the evaluative decision-making involved in assessing efficiency. Framings of efficiency relatively prioritise the ends and interests, and reflect the purposes, of particular individuals and agencies, whether this is explicitly recognised or not. This raises distinctively ethical questions. Are any agents unfairly burdened by the way that efficiency is characterised? Do the purposes that drive the framing of efficiency reflect what matters? What values are (implicitly or explicitly) endorsed in the framing of efficiency and are they the values that the service or system should embody? Answering such questions is not an empirical matter. Nor is it a matter of scrutinising the concept of efficiency itself. Rather, giving considered answers to these questions requires context-specific deliberation about what the system and service in question is and does and what it ought to be and do. This is not exhausted by extensive stakeholder engagement with

those affected by services, nor by modelling alternative efficiency-oriented interventions to estimate their expected effects, including their impact on different stakeholders—although both of these might usefully form a part of a deliberative process. Deliberative decision-making about efficiency will require consideration of such information as part of a reflective, evaluative and open-ended process, which considers the different ways of interpreting the available evidence and is genuinely open to different outcomes.

Characterising efficiency and managing the trade-offs that this involves entails attending to the different perspectives that individuals and institutions have on purposes of a service and the resources available to it. Acquiring knowledge of these perspectives is also a normative and contested matter. It is crucial to recognise that understanding system functioning for the purposes of characterising efficiency is an activity which can be conducted in an indefinite number of ways, each achieving a different understanding and capturing different perspectives. For example, the qualitative stakeholder engagement that WORKTECC undertook included semi-structured interviews, workshops and some shadowing of staff, and these yielded different insights into how particular actors experience and conceptualise the service. Asking people to describe the actual and ideal functioning of a system is not the only way of coming to understand how that system functions or ought to function. Observation of the system in process—for example, observing a manager plan the operations of a service, or observing a home visit—will paint a different picture. Surveying or observing different individuals, in different roles, could also inform understanding. Consideration of recorded documentation describing system function—including statements of purpose, risk and incident logs, past schedules and scheduling tools or guidance, and also data about travel time, staffing, the patient population, and clinical outcomes—would show something different again. These alternative accounts might well overlap in many respects with the account of the service as declared and described by stakeholders, but they may also diverge. There are a range of ways to acquire knowledge about a system for the purpose of mapping and modelling efficiency and the management of different sources of information presents modellers with a series of choices. These are not only choices about who and what to ask, but also choices about how to manage conflicts, how to understand considerations that are only recognised as significant from certain perspectives, and how to identify and characterise ‘good’ functioning as distinct from past or current functioning.

Making judgements about the allocative efficiency of a system is thus highly dependent on the process of coming to understand the different purposes and components of a system and how they fit together, and the perspectives reflected in this process. However, in practice, anyone who wants to measure and assess efficiency must make difficult choices about how to characterise the health system in question and so, implicitly, whose interests to include. Not only is it not practically possible to model all resources and ends, but resource-constrained systems must prioritise some ends over others when they cannot pursue all desired aims at once. In the next section, we explore the relationship between models of efficiency and the systems that they describe and suggest that such models are not just passive descriptions, but rather active, creative phenomena.

4. THE MAP IS NOT THE TERRITORY

Any model of efficiency will end up excluding or substituting certain factors. This could be for a variety of reasons. Some inputs and outputs are very difficult to characterise definitely or measure directly—personal and interpersonal goods, such as emotions, well-being, communication, and trust plausibly fall into this category. Even if a service cannot fulfil its function without them, it may be very difficult to include them directly in a model. Models can use proxy measures, which are more easily identifiable and measurable, but these indirect measures will only capture their objects in an imperfect way. Staff well-being, for instance, may be indirectly represented by a number of factors—hours worked, hours of overtime worked, use of skills, and so on—and a model may endeavour to place limits or restrictions on some of these factors in an effort to ensure that constructions of efficiency do not ignore well-being. But this will by no means guarantee that staff well-being is adequately protected or secured as a system outcome.

Other inputs and outputs will be excluded from models because their political or sensitive nature makes their inclusion inappropriate. There may, for instance, be good reason for the care of a particularly complex and challenging patient or client to be shared, rather than for them to be seen by the same staff member week after week, even though this might conflict with an aspiration to provide continuity of care or assistance. Or, if two staff members do not work well together due to a personal conflict, there might be good reason not to schedule them for shared jobs, even when this creates other scheduling issues. These might well be considerations that a sensitive manager would take into account when planning shifts and they would be sensible to do so—such factors will very plausibly impact on operational and clinical outcomes, as well as personal ones. But including these considerations in a model would require them to be articulated and formalised as institutional facts in a way that at best would be inappropriate and at worst would compound the problems in question and represent professional malpractice. Explicitly including factors like this in a formal model has the potential to crystallise them in a way that can be avoided by operating with discretion in less formalised decision-making.

So there will, inevitably, be a gap between the actual functioning (and indeed the *good* or *ideal* functioning) of a system or service and a model of it. The map is not the territory. This is not in itself a problem and indeed is part of the nature of models, which necessarily involve abstraction and simplification. But modelling efficiency is not just a descriptive task—such modelling is intended to be used to help change the management of resources in order to improve service efficiency. And prescriptively using partial models has the potential to introduce problems. If a model misses out or misrepresents key aspects of system functioning, its recommendations will fail to adequately reflect efficiency gains. But if the model is understood or presented as exhaustively characterising efficiency, then following the model's suggestions may generate inefficiencies that are difficult to characterise as such. Furthermore, using such a model has the potential to construct a service in the image of a partial model and so implicitly deprioritise those elements that are excluded. When those elements are in fact important to system functioning, this has the potential to result in inefficiencies from an institution or system perspective. But it also has the potential to be extremely costly to individuals in relation to their personal well-being and emotional resources. It is crucial to recognise that a model of a complex system that looks neat and tidy—that breaks it down into

a finite set of clearly defined inputs and outputs, which inter-relate in finite and clearly defined ways—need not, indeed likely will not, reflect the actual functioning of the system (Jacobs 2016). Nor need it reflect the *good* functioning of the system in realistic conditions.

There are a number of ways of avoiding these pitfalls. The models developed as part of the WORKTECC project indicate some possible ways forward. First and foremost, WORKTECC embraces a range of relevant indicators and perspectives, acknowledging that modelling efficiency is a complex, multi-dimensional matter. The broad, qualitative stakeholder engagement sought to identify and understand different perspectives on the actual and ideal function of a home healthcare service. But, within the context of this complexity and plurality, WORKTECC offers technical assessments of diverse aspects of efficiency such as staff costs (disaggregated by salaried staff and agency staff), workload balance within and between teams, the proportion of time that staff spend on activities that could be performed by a lesser-skilled member of staff, and the average number of different individuals that visit patients. Unless efficiency is operationalised via these more technical assessments, it is difficult to see how practical improvement efforts could get off the ground. In practice, given the resource-limitations of healthcare organisations, difficult decisions must be taken about how to manage scarce resources and which set of ends to prioritise, given that it is likely to be impossible to do everything that is valued and desired. However, using technical assessments of efficiency alone to make such decisions risks false economy, if it fails to factor into decision-making a series of variables which are in fact crucial for good system functioning. Failure to take account of these factors could mean that what is claimed to be more efficient ends up being no such thing.

The WORKTECC approach takes steps to avoid the necessary reductiveness of models leading to a reductive conceptualisation of efficiency more broadly. It does this, in part, by consciously mapping the complementarities and tensions between different perspectives—that is, it indicates how operational changes would impact on different system outcomes in relation to different stakeholders, rather than producing a single judgement of what would be most efficient. By focusing on the impact of particular efficiency-oriented interventions, and by highlighting the trade-offs involved, such an approach avoids making definitive claims about efficiency but expressly supports deliberation about the potential impact and feasibility of alternative possible efficiency-oriented interventions. Efficiency gains are thus presented as normative choices, incorporating trade-offs, rather than unequivocal improvements. In short, the WORKTECC modelling frameworks are explicitly characterised as partial and suggestive and as generating useful tools rather than comprehensive solutions. Models are presented as tools that sit within, and form only one part of, a decision-making process, rather than as alternatives to a decision-making process. Assessments of technical efficiency can contribute to more reflective and deliberative decision-making about allocative efficiency, but they shouldn't on their own be taken to give definitive and decisive answers about the efficiency of health systems.

But no matter how sensitive researchers who develop models are to the potential dangers of overly reductive use of models in decision-making, there remains a risk that their caveats are, in practice, ignored. This raises a broader set of ethical concerns about creating and using models in healthcare contexts. In order to use constructions of efficiency as practical

decision-making tools, it seems inevitable that they must be operationalised in relatively limited, technical ways. But there is an associated risk that limited, technical characterisations are invoked as if they are exhaustive of efficiency—even if they are designed with an intent that they are used as part of deliberative decision-making. The responsibility of managing this risk cannot fall solely to modellers. We suggest that it is, in particular, a responsibility of managers and those with operational decision-making power within institutions to remain alert to the limitations of models as decision-making tools.

5. CONCLUDING THOUGHTS

We have tried to emphasise the importance of taking seriously and holding together both more technical and more evaluative approaches to efficiency. Operationalising efficiency generates useful tools but only if used with awareness of the major limitations and framing effects of operationalisation. Using more open-ended lenses to think about efficiency helpfully highlights innumerable uncertainties and contestations but also risks impracticality. The challenge is to combine and consciously balance reductionism and indeterminacy.

We have used the WORKTECC example to investigate this domain and we draw on it to conclude by making three linked points—two substantive and one procedural—that summarise our argument. First, modelling efficiency necessitates narrowing in on some specific concerns and excluding others. In the case of home healthcare services, for example, an emphasis on optimising levels of service activity risks ‘framing out’ aspects of the lifeworld of both service users and staff. The danger here is that important elements of the perceptions, preferences and quality of experience and life of service users and staff may be neglected by constructions of efficiency that focus on immediate processes rather than broader service effects. Furthermore, we have suggested that if these limited constructions enter in an unreflective way into the efficiency discourses in circulation, then this form of neglect may be reified, reproduced and magnified. Of course, personal and emotional concerns may, at least in principle, be captured by a focus on other values but, as we have shown, they are equally relevant to thinking about efficiency. There is no easy solution here. Foregrounding more of these concerns in a model of efficiency would, in turn, be at the expense of other emphases.

Second, and more broadly, for a model to be of use to a service it is likely that it will rely on a particular—currently dominant—reading of service function and system context. This is because, as we have underlined, conceptions of the efficiency of a health service are always framed—they focus on specific inputs and outputs, seen from particular perspectives and against assumptions about system purposes. Not only can alternative efficiency framings be imagined but, in many cases, they will already be emerging in the policy environment. For example, should a home healthcare service be judged against a broader healthcare prevention frame or perhaps a more ‘pastoral’ social care frame? As well as meeting immediate healthcare treatment needs, how far should the success of such a service be judged by indicators of longer-term ends, whether that be keeping people out of hospital or helping to provide, and connect people with, support to live socially connected lives that they value? This is an area where decisions about salient efficiency frames are likely to be highly volatile and can easily shift as services are reformed, systems reconfigured and budgets re-allocated such that authoritative constructions of purposes, and the balance of purposes, are redrawn. These instabilities and

contestations will not only reflect evolving conceptions of particular services but parallel kinds of evolution in the funding, design and affordances in the broader landscape of formal and informal provision. Even though these longer-term concerns are extremely difficult to quantitatively model, they play an important qualitative role in constructions of models and measures of service efficiency. Home healthcare services provide a potent example because they sit in the middle of a highly contested policy arena that contains lively debates about purposes and the costs and benefits of service co-ordination and integration. But this argument—about how conceptions of service quality are made unstable by system or policy change—is of widespread relevance.

Third, and finally, the modelling approach taken by WORKTECC highlights one way of holding together the technical and evaluative dimensions of efficiency and supporting the crucial dialogue between the two. In this case it involves technical modelling of constructions of efficiency in a way that embodies careful attention to the complexities of so doing. This includes explicitly acknowledging the need for trade-offs between different values and perspectives and presenting resulting models not as single comprehensive solutions but as a range of tools to be deployed within broader deliberations about system performance and reform. There are, no doubt, many other ways to pay due respect to broader as well as narrower conceptions of efficiency, but these will all arguably share the imperative not to ‘close off’ opportunities for seeing efficiency as a deliberative and evaluative concern and not merely a technical one. This, we suggest, is not just an imperative for modellers and service decision-makers but for anyone participating in, or acting as a critical consumer of, discussions about efficiency.

We have suggested that efficiency is different from other values but not perhaps in the way that might have been assumed. As indicated at the start of the paper, the idea that a service is efficient may seem to be a relatively ‘hard’ one when compared to other ideas—such as being responsive or equitable—which appear ‘soft’ and intangible by comparison. But, as we have highlighted and illustrated, real world assessments of efficiency are unavoidably value-laden because they must draw on particular, contestable characterisations of the purpose and function of systems. Efficiency only has a credible claim to something like value-neutrality when considered at the most abstract or formal level—the level at which it is largely empty.

Nonetheless whilst empirical claims about efficiency depend on claims about other values, they are not reducible to them. Efficiency is an important value in its own right. It focuses attention on the challenge of using resources to their best effect. In relation to public services this may, for specific purposes and to a limited extent, be constructed as about maximising a particular output (per unit of input). But it is, more generally, about ‘making ends meet,’ that is, seeking to identify optimal combinations of ends that can be met by the combination of available (or attainable) resources. Discourses and practices of efficiency are, therefore, always bound up with debates about other service values and purposes. Deliberation and decision-making about efficiency cannot be replaced with models and measures because evaluative processes are inherent to the construction and use of measures and models. In this sense ‘doing efficiency’ is always also ‘doing ethics.’

REFERENCES

Alexander, J. K. 2008. *The mantra of efficiency: from waterwheel to social control*. Baltimore: Johns Hopkins University Press.

Alexander, J. K. 2009. The Concept of Efficiency: An Historical Analysis. *Philosophy of Technology and Engineering Sciences*, Elsevier: 1007-1030.

Cantor, V. J. M. and K. L. Poh. 2017. "Integrated Analysis of Healthcare Efficiency: A Systematic Review." *Journal of Medical Systems* 42(1): 8.

Culyer, T. 2015. *Efficiency, Equity and Equality in Health and Health Care*. University of York, Centre for Health Economics.

Cylus, J., I. Papanicolas, P. C. Smith and World Health Organization. 2016. *Health System Efficiency: How to Make Measurement Matter for Policy and Management*, World Health Organization, Regional Office for Europe.

Davies, G. 2022. "Efficiency Savings Require Learning Past Lessons." National Audit Office Press Release. December 8. Accessed 13 December 2023. <https://www.nao.org.uk/press-releases/efficiency-savings-require-learning-past-lessons/>.

Fogel, R. W. and S. L. Engerman. 1995. *Time On The Cross: The Economics Of American Negro Slavery*. New York, W.W. Norton.

Grieco, L., M. Utley and S. Crowe. 2021. "Operational Research Applied to Decisions in Home Health Care: A Systematic Literature Review." *Journal of the Operational Research Society* 72(9): 1960-1991.

Gutman, H. G. 1975. *Slavery and the Numbers Game: A Critique of Time on the Cross*. Urbana, University of Illinois Press.

Hilt, E. 2020. "Revisiting Time on the Cross After 45 Years: The Slavery Debates and the New Economic History." *Capitalism* 1(2): 456-483.

Hussey, P. S., H. De Vries, J. Romley, et al. 2009. "A Systematic Review of Health Care Efficiency Measures." *Health Services Research* 44(3): 784-805.

Institute of Medicine (US) Committee on Quality of Health Care in America. 2001. *Crossing the Quality Chasm: A New Health System for the 21st Century*. Washington, D.C., National Academy Press.

Jacobs, J. 2016. *The Death and Life of Great American Cities*, Vintage.

Knight, J. 1992. *Institutions and Social Conflict*, Cambridge University Press.

Le Grand, J. 1990. "Equity Versus Efficiency: The Elusive Trade-Off." *Ethics* 100(3): 554-568.

Maxwell, R. J. 1984. "Quality Assessment in Health." *British Medical Journal (Clinical research ed.)* 288(6428): 1470.

McGlynn, E. A. and P. G. Shekelle. 2008. *Identifying, Categorizing, and Evaluating Health Care Efficiency Measures*, Agency for Healthcare Research and Quality.

Merry, L. and S. Gainsbury. 2023. "What Was Austerity's Toll on the NHS Before the Pandemic?" Nuffield Trust Blog Post. June 21. Accessed 13 December 2023. <https://www.nuffieldtrust.org.uk/news-item/what-was-austeritys-toll-on-the-nhs-before-the-pandemic>.

Mortimer, F., J. Isherwood, A. Wilkinson and E. Vaux. 2018. "Sustainability in Quality Improvement: Redefining Value." *Future Healthcare Journal* 5(2): 88-93.

Palmer, S. and D. J. Torgerson. 1999. "Definitions of Efficiency." *BMJ* 318(7191): 1136.

Pereira Gray, D. J., K. Sidaway-Lee, E. White, A. Thorne and P. H. Evans. 2018. "Continuity of Care with Doctors—A Matter of Life and Death? A Systematic Review of Continuity of Care and Mortality." *BMJ Open* 8(6): e021161.

Santana, M.-J., S. Ahmed, D. Lorenzetti, et al. 2019. "Measuring Patient-Centred System Performance: A Scoping Review of Patient-Centred Care Quality Indicators." *BMJ Open* 9(1): e023596.

Sassi, F., J. Le Grand and L. Archard. 2001. Equity Versus Efficiency: A Dilemma for the NHS: If the NHS Is Serious About Equity It Must Offer Guidance When Principles Conflict, *BMJ* 323(7316): 762-763.

Universities UK Efficiency and Modernisation Task Group. 2011. Efficiency And Effectiveness in Higher Education: A Report by the Universities UK Efficiency and Modernisation Task Group. London, Universities UK.

Van de Poel, I. 2009. Values in Engineering Design. *Philosophy of Technology and Engineering Sciences*, Elsevier: 973-1006.

Varabyova, Y. and J.-M. Müller. 2016. "The Efficiency of Health Care Production in OECD Countries: A Systematic Review and Meta-Analysis of Cross-Country Comparisons." *Health Policy* 120(3): 252-263.

Williams, A. 1988. "Priority Setting in Public and Private Health Care: A Guide Through the Ideological Jungle." *Journal of Health Economics* 7(2): 173-183.

Zakowska, I. and M. Godycki-Cwirko. 2019. "Data Envelopment Analysis Applications in Primary Health Care: A Systematic Review." *Family Practice* 37(2): 147-153.