



Efficient Inefficiency: Organisational challenges of realising economic gains from AI

Stuart Mills^{a,b,*}, David A. Spencer^c

^a Economics, University of Leeds, UK

^b Visiting Fellow of Behavioural Science, London School of Economics and Political Science, UK

^c Economics and Political Economy, University of Leeds, UK

ARTICLE INFO

Keywords:

Artificial Intelligence
Bounded Rationality
Organisational Efficiency
Productivity
Superfluous Work

ABSTRACT

Organisations are increasingly using artificial intelligence (AI). Where AI performs productive tasks more efficiently than humans, organisations will benefit economically through increases in productivity. However, if AI is deployed to undertake unproductive, superfluous tasks, the efficiency benefits will be reduced, even if these tasks are performed more efficiently than a human could, because the said tasks are inefficient to begin with. We call this eventuality ‘efficient inefficiency.’

We outline several reasons why superfluous tasks are created by managers and why they persist in organisations, drawing on an array of behavioural, managerial, and sociological literature. We argue bounded rationality accounts for why managers often fail to identify superfluous tasks, coupled with organisational conflicts which often incentivise their creation. These factors impede the ability of organisations to avoid efficient inefficiency. Restructuring organisations to promote knowledge sharing and align stakeholder incentives may reduce, though not eliminate, the risk of efficient inefficiency.

1. Introduction

In organisations, managers must decide which technologies are used, and how. Broadly, managers in organisations use new technologies to increase efficiency. In conventional capitalist organisations, the aim remains to ensure maximum return for the smallest cost outlay. Technology is used to maximise output from workers while reducing production and other costs. The potential and actual efficiency benefits of technology remain an important focus within management studies (Audretsch and Belitski, 2023; Dodgson, 1991) and these benefits have gained new attention following recent developments in artificial intelligence (AI).

AI is already entering the workplace. ChatGPT, developed by OpenAI, is the fastest growing online service in history (Ahuja, 2023). Microsoft is integrating AI analytics and co-piloting tools into its suite of Office software (McGee and Murgia, 2023). Amazon has patented and deployed various AI-powered technologies across its business activities (Delfanti and Frey, 2021). The professional software company SAP has recently revealed a ‘roadmap’ aimed at incorporating generative AI into its products (SAP, 2023). Other new AI technologies will likely be developed and adopted in the future.

The aim to raise efficiency is central to economic perspectives on AI (Acemoglu and Johnson, 2023). These perspectives, along with more popular discussion in the media, link AI adoption with the opportunity for huge efficiency gains (Acemoglu, 2024; Brynjolfsson and McAfee, 2014). Yet, such discussion often focuses on the *technical* potential for AI to perform existing workplace tasks (Frey and Osborne, 2017). Discussions of *what* tasks are performed, *why* tasks are performed, and the organisational challenges of determining *how* technologies should be used, are often overlooked. For instance, some emerging studies suggest recent AI innovations may produce economic benefits through productivity gains (e.g., Cambon et al., 2023; Dell’Acqua et al., 2023). However, these studies are often isolated experiments, lacking an organisational context, and focused on simple productivity measures.

In this paper, we argue that the efficiency gains of AI technologies cannot be removed from an organisational context, and the conflicts contained therein. From an efficiency perspective, the technical question ‘can a particular technology perform a task?’ is less important than the organisational question ‘should the task be performed?’ and the managerial question ‘who should perform the task?’ These questions provide the key research focus for this paper. We are aware that the issues we address over technology adoption and its relationship to

* Corresponding author.

E-mail address: s.mills1@leeds.ac.uk (S. Mills).

efficiency are general – in effect, they potentially apply to all instances where technology is used for efficiency improvement. For the purposes of this paper, however, we focus on the example of AI, given both its topicality and the somewhat unique questions it raises for organisations.

We challenge the often-accepted view that AI will produce higher organisational efficiency. Rather, we argue AI will only realise economic benefits for organisations if the tasks it is directed towards are productive in the first instance. If these tasks are not productive, but *superfluous*, AI will contribute limited economic value. For instance, AI could be deployed more productively elsewhere in organisations for little or no additional cost. Efficiency gains from the adoption of AI, when applied to superfluous tasks, are only *relative* economic gains for organisations, not absolute ones. In making this point, we originate the concept of ‘efficient inefficiency’. Defining this concept and showing how it can be applied to understand the AI adoption decision is the central contribution of this paper.

Efficient inefficiency rests on organisations creating and sustaining superfluous tasks – namely tasks that do not contribute to the efficiency of organisations. Drawing on different sociological theories of ‘bullshit jobs’ and bureaucracy (Graeber, 2018, 2015; Alvesson and Spicer, 2012; Spicer, 2013; Tiratelli, 2022; Weber, 2019), cybernetic management theories of organisational control (Beer, 1993, 1979; Medina, 2014), and behavioural organisation theories of bounded rationality (Cyert and March 1963; Egidi, 1997; Simon, 1991, 1947), we argue managers often face various incentives to create superfluous tasks.

For illustrative purposes, we distinguish between senior managers (who take and impose decisions about technological adoption) and more junior managers (who receive and (in principle) follow instructions about technological adoption). This distinction allows us to demonstrate how superfluous tasks can emerge through managerial interests that differ depending on power and status within organisational hierarchies. We propose that this same division of management enables superfluous tasks to be sustained in the long run. The division between senior managers, who decide, and junior managers, who respond, creates a knowledge problem between these two groups. This problem is characterised by bounded rationality (Simon, 1997, 1947).

Senior managers must determine how technology is used, but in complex organisations, they face the problem of assimilating too much information. Furthermore, they face uncertainty in their decisions, as they do not (and often cannot) know how junior managers (and also workers) will respond to their decisions. This makes it difficult for senior managers to determine superfluous tasks from productive ones. Junior managers do not (and potentially cannot) know why senior managers are taking particular decisions. This creates incentives for junior managers to increase uncertainty as a defensive strategy, including creating new superfluous tasks and not disclosing existing ones. These incentives mean that technologies – including AI – can be used to support superfluous tasks, resulting in efficient inefficiency.

We identify efficient inefficiency primarily as an organisational and managerial challenge, not a technological one. We explore the implications of efficient inefficiency for practitioners and consider strategies which might be adopted to ameliorate it. One strategy may be to restructure organisations, reducing the factors that create the knowledge problem we describe. Senior managers may give junior managers greater autonomy and decision-making power, spreading the ‘cognitive load’ of the technology adoption decision. Giving junior managers more autonomy and decision-making power, coupled with protections to their status and job security, would also align junior managers’ incentives with those of senior managers, reducing organisational uncertainty.

Drawing on perspectives that examine the notion of holacracy (Bernstein et al., 2016; Laloux, 2014; Mintzberg, 1979; Robertson, 2015; Schell and Bischof, 2021), we suggest organisations with flatter hierarchies, aligned stakeholder interests, and democratic decision-making may be better equipped to tackle the root causes of efficient inefficiency and thus may be better positioned to realise the economic benefits of technologies such as AI. Interestingly, *unlike* other

technologies, we suggest AI may have advantages in supporting such restructuring of organisations as a tool for augmenting senior manager decision-making (Simon, 1987a, 1987b), facilitating democratic decision-making mechanisms (Specian, 2024, 2023) and promoting organisational innovation (Bouschery et al., 2023; Broekhuizen et al., 2023; Medina, 2014).

We also consider the theoretical implications of efficient inefficiency and the opportunities the concept raises for several disciplines. In particular, we offer a novel perspective on the ‘Solow paradox’ found in the economics of technology literature. This literature has long been concerned with the lag between technology adoption and efficiency benefits. We suggest efficient inefficiency is an important explanation for this lag.

The paper is organised as follows. The Theoretical Perspectives section reviews the literature on the nature and origins of superfluous tasks. The Bounded Rationality and AI Adoption section explains why these tasks persist and why their persistence frustrates the effective deployment of technology, leading to efficient inefficiency. The Implications for Practitioners section considers various strategies to support practitioners, while the Implications for Theory section discusses the link between efficient inefficiency and the Solow paradox. The Future Perspectives section considers how efficient inefficient may be further developed by business researchers and practitioners. The main conclusions follow.

2. Theoretical perspectives on superfluous tasks

If a technology can be used to perform a task more efficiently, it may realise economic benefits for an organisation by reducing the costs associated with the task and unlocking organisational resources to exploit productive opportunities elsewhere (Autor, 2015). Yet, these economic benefits are largely independent of the technology. Instead, they overwhelmingly depend on the task to which the technology is applied. An organisation may only realise absolute economic gains if a task is economically worthwhile and technology can improve the efficiency of this task’s performance, either through automating the task or through augmenting the task to support manual performance (Acemoglu and Restrepo, 2019).

However, if the task serves little or no productive purpose, which is to say, it is superfluous, then applying technology will not realise absolute economic gains. It may realise relative gains, as the superfluous task is performed more efficiently. But overall, technology will not be deployed efficiently, and any increase in efficiency must be measured from an unproductive baseline. Performing superfluous tasks more efficiently represents a *drag* on efficiency – efficiency may increase, but by less than if the superfluous task was simply eliminated. This outcome, broadly, is what we call efficient inefficiency.

Efficient inefficiency only matters – and indeed, only exists – if one expects superfluous tasks to exist within organisations. Neoclassical economic theory would suggest that organisations replete with these tasks would, in the long run, be unable to survive within competitive markets (Alchian, 1950; Friedman, 1953; Simon, 1981). This is because more efficient rivals would consistently out-compete these organisations, resulting in their eventual failure. Only efficient organisations are said to prevail and survive in the long run.

Yet, this argument, while seemingly compelling, ultimately does little to undermine the *possibility* of superfluous tasks persisting in organisations. If the conditions by which superfluous tasks are created and sustained are found in all (or, at least most) organisations, superfluous tasks could persist while organisations remain competitive with one another (Egidi, 1997). A competitive advantage would be afforded to those with *relatively fewer* superfluous tasks, but there is no particular reason to believe that, in the long run, organisations would eliminate all superfluous tasks. Indeed, taking a transaction costs perspective (Williamson, 1981), if the costs of identifying and eliminating these tasks (e.g., search and monitoring costs) are greater than the costs of

continuing with them, even the most competitive organisations would continue to undertake superfluous work. Competition, in short, need not eliminate superfluous tasks and thus efficient inefficiency.

Various literatures describe how superfluous tasks are created within organisations. Most eye-catchingly, the recent bullshit jobs perspective argues that an assortment of sociological factors contribute to the creation of work which serves no obvious productive purpose (Graeber, 2018; Spicer, 2013). Managers in organisations, it is argued, will deliberately create superfluous tasks because they gain personal benefits from doing so. Some cybernetic management perspectives (Beer, 1993, 1979; Medina, 2014) offer additional arguments for the emergence of superfluous tasks as means of control, as do various sociological accounts of the emergence of bureaucracy in organisations (Graeber, 2015; Tiratelli, 2022; Weber, 2019) and behavioural accounts of the role of bounded rationality (Cyert and March 1963; Egidi, 1997; Simon, 1997, 1947).

For the purposes of the discussion below, which focuses on the elaboration of efficient inefficiency, we distinguish between two sets of managers that exist within a typical capitalist organisation. Firstly, we identify senior managers who command junior managers and make decisions about how technology is deployed, and how junior managers use it. Secondly, we identify junior managers, who, in principle, follow the instructions of senior managers, but themselves have authority over workers within an organisation. They also have some autonomy over their own behaviour. It is useful to distinguish between these two types of managers, as it allows for a more realistic representation of managerial roles and interests within a typical capitalist organisation (e.g., Cyert and March 1963), particularly in regard to how superfluous tasks emerge and persist, and how technology is adopted (Dodgson, 1991; Faraj and Pachidi, 2021). Though, we recognise that such a hierarchical structure may not be representative of all organisations, and indeed, in Section 4, we consider how alternative structures (such as holacracy) may resolve some of the difficulties associated with efficient inefficiency.

The concept of bullshit jobs derives from various survey data suggesting a significant minority of people in Western capitalist countries do not believe their jobs contribute productively to the organisations in which they work, or to society at-large (Graeber, 2018). The concept has been further linked to classical institutional economic accounts of organisations based on tensions between organisational goals and individual motivations (Dean, Dadzie and Pham, 2022). Bullshit jobs are said to arise for various reasons, particularly the exercise and display of power by senior managers.

Research has shown how senior managers enjoy exercising power via the recruitment of extra staff. The hiring of so-called ‘flunkies’ and other superfluous staff can signal status within an organisation, which senior managers particularly value (Graeber, 2018). Some have also suggested that bullshit jobs may be created as a defensive strategy employed by junior managers against senior managers. Junior managers may wish to invent work – either for themselves or their staff – in order to appear more important and indispensable (Alvesson and Spicer, 2012; Spicer, 2013). Junior managers might also create bullshit jobs because they lack authority within the organisation. So-called ‘duct-tapers’ are said to temporarily fix problems which could be permanently solved, but for various reasons (e.g., organisational ignorance, myopia, a lack of funds and authority to implement a solution) go unsolved (Graeber, 2018). Duct-tapers exist partly because the organisational structure prevents the implementation of systemic solutions.

Some cybernetic management theorists have also argued that senior managers prioritise control of organisations over efficiency (Beer, 1993, 1979; Medina, 2014). Senior managers may place demands on those below them to produce particular documents, follow particular processes, and work particular times, not for clear economic reasons, but instead for the purpose of maintaining control of the organisation (Marglin, 1974), and signalling their own position and authority in the organisation (in a manner analogous to the employment of ‘flunkies’).

For instance, one study of managers at a software company noted that particularly controlling managers would penalise staff for not working in specified ways, or at specified times, despite those staff having completed all work required of them (Perlow, 1998). This suggests that an important reason for these working requirements was managerial motives of control, rather than necessarily organisational productivity.

Related to cybernetic theories of control are some behavioural perspectives on organisations and control (Cyert and March 1963; Egidi, 1997). These conceive of firms as adaptive entities made up of various stakeholders (e.g., owners, senior managers, junior managers and workers). These stakeholders, at any given time, may have conflicting interests. For instance, a junior manager may want status, a worker may want a pay rise, and a shareholder may want maximal profit. Conflicts between the interests of these stakeholders can create the basis for inefficiency, not least through the deployment of resources to combat the conflicts themselves.

Because stakeholders are said to be boundedly rational (Simon, 1955), influential decision-makers (e.g., senior managers) do not and cannot tackle all conflicting interests at once. Rather, they must focus on the most serious conflicts. For instance, pressures arising from shareholders, or industrial disputes coming from workers’ unions. Less serious conflicts, such as those arising from junior managers’ desire for status, are not given substantial attention, and inefficiencies stemming from these conflicts are permitted (leading to ‘organisational slack’) if only tacitly so that attention may be focused elsewhere (Cyert and March 1963).

Acting in this way enables senior managers to maintain a degree of control within the organisation given their scarce (cognitive) resources (Cyert and March 1963; Hodgkinson et al., 2023), but permits junior managers to engage in superfluous activities which may not benefit the organisation. For instance, a recent study found that junior managers frequently manipulated how long their work took, and the apparent difficulty of their work, so as to retain control over their work time, rather than ceding that time to the direction of more senior managers (Lupu and Rokka, 2022).

Some have argued that the bounded rationality of managers allows superfluous tasks to persist in the long run, and that the number of superfluous tasks should increase as organisations grow in size and complexity (Egidi, 1997). For instance, one study of Japanese firms found that organisational size is a significant predictor of the allocation of organisational resources to tasks which do not benefit the firm (Karube et al., 2009).

The above arguments are not dissimilar to some sociological theories of bureaucracy (Tiratelli, 2022). Managers may create bureaucratic structures not just to monitor workers, but also to receive data from which decisions can be made. When bureaucracy originates productive insights, it may create an advantage for some senior managers and the organisation in which they work. Though, as organisations become more complex, bureaucracies must expand, and over time, their maintenance can become an end in itself (Weber, 2019). For instance, recent studies of bureaucratic structures have proposed ideas such as ‘functional stupidity’ and ‘absurdocracy’ to describe instances of workers engaging in activities which they (and their managers) know to be pointless, but which are necessary to satisfy administrative systems which neither they nor their managers can change (Alvesson and Spicer, 2016; Alvesson and Stephens, 2024; Buchanan and Badham, 2020).

As with some of the older economic perspectives discussed above, it has been argued that senior managers – who are often in control of bureaucratic structures – may seek to extend their bureaucratic reach for the purposes of displaying and maintaining their power, organisational knowledge, and status (Graeber, 2015; Tiratelli, 2022). Sociological theories, then, offer antecedents for modern accounts of bullshit jobs, where superfluous tasks are seen as endemic in workplaces.

The benefit of considering cybernetic, behavioural, and sociological theories in conjunction with the bullshit jobs thesis is that it provides further supporting context for the existence of organisational

inefficiency. The bullshit jobs thesis has been criticised for its emphasis on individual, subjective perceptions of work (Tiratelli, 2022; Spencer, 2022). For instance, an environmentalist would probably consider their job to be bullshit, if they worked for an oil or mining company, because the organisation does not align with their individual values and beliefs. Such subjective accounts, however, say nothing about the actual productive value of their job nor the value of their job on its own terms (e.g., it may be highly paid and offer positive intrinsic rewards). That said, the cybernetic, behavioural, and sociological perspectives confirm a core element of truth in the bullshit jobs thesis, namely that superfluous work is not merely (or even largely) an artefact of subjective experiences (Dean, Dadzie and Pham, 2022), but the product of organisational structures and managerial decisions (Cyert and March 1963; Tiratelli, 2022). Taken together, all are useful in showing how and why superfluous tasks arise and persist within organisations.

3. Bounded rationality and AI adoption

Efficient inefficiency arises when technology is used to perform superfluous tasks more efficiently. However, even accepting that superfluous tasks exist within organisations, efficient inefficiency need not arise if senior managers can identify these tasks before deciding how technologies such as AI are used. Therefore, an outstanding question is why might senior managers be unable to identify superfluous tasks (beyond the fact that they, themselves, may also create such tasks)? One compelling answer comes from the presence of bounded rationality, and more specifically, the conflicting goals of senior and junior managers.

Bounded rationality is a well-known theory of human decision-making (Simon, 1997, 1981, 1955), with applications in organisational behaviour (Cyert and March 1963; Dosi and Marengo, 2007; Egidi, 1997; Schoemaker, 1993; Simon, 1997, 1947) and management studies (Hodgkinson et al., 2023; Stubbart, 1989). It argues that people must make decisions under informational and cognitive limits. This may be linked to the complexity of the decision being made, which limits how information can be cognitively integrated into the decision-making process. Another form is information uncertainty. Uncertainty arises because there is information that a person cannot know, either by design (e.g., people are not incentivised to disclose information) or by nature (e.g., information concerns phenomena which have not yet happened).

Boundedly rational decision-makers engage in adaptive behaviours to overcome these challenges (Simon, 1981). These may include simplifying decisions and using mental shortcuts or heuristics (Simon, 1991). Actions taken are rarely 'utility maximising,' as neoclassical economics generally assumes. Rather, they entail 'satisficing' or 'good enough' actions that allow a person to respond to and resolve whatever decision is facing them, while still making an acceptable choice given the limitations surrounding the decision (Simon, 1955, 1947).

Senior managers are boundedly rational when they come to make decisions concerning technology adoption. They must synthesise large amounts of organisational information into their decisions. The decision to deploy a technology such as AI to perform a given task requires them to know how that task relates to all other tasks undertaken by others in the organisation. Without this knowledge, they may not be able to determine if the task is productive or superfluous. Where a senior manager must make such decisions for multiple different teams, the cognitive challenge of integrating all relevant information increases (Simon, 1981, 1947). Given the presence of bounded rationality, senior managers are therefore unlikely to make optimal decisions (i.e., those which correctly identify and eliminate superfluous tasks), and are more likely to 'satisfice' (i.e., some tasks to which a technology is deployed will be superfluous).

An interesting example of how such satisficing leads to efficient inefficiency comes from Simon's (1981) discussion of telegram handling in a US diplomatic office during the 1960s, where he worked as an administrative consultant. During international incidents, the office would be inundated with telegrams, which were printed using

teleprinters. With a limited printing speed, telegrams would soon back up, and vital information would be delayed in reaching diplomats. Managers intuitively understood this challenge as a technical one resulting from a lack of printers and responded in kind by purchasing more. Nevertheless, as Simon (1981) reports, the problem of delayed information continued, even with the additional printers.

For Simon (1981), the office's real problem was printing information that diplomats were *never going to use*. Adding more printers simply allowed more superfluous information to be printed and did not enhance the actual productivity of the office. Eliminating superfluous printing would have been a significant challenge, as it arose from the uncertainty of international incidents, and thus the bounded rationality of diplomats. Yet, the satisficing 'solution' of adding more printers ultimately did not solve the problem, and in some ways *exacerbated* it by enabling *more* printing – it simply added to the volume of information that diplomats had to assimilate and understand. This is an exemplar of efficient inefficiency. For Simon (1981), only solutions which responded to the bounded rationality of diplomats could have solved the delay problem.

Managers who are forced to satisfice may take decisions based on immediately apparent benefits rather than examining the key, underlying drivers of (or impediments to) productivity. As mentioned in the above example, the decision by office managers to purchase more printers was likely the result of diplomats satisficing to increase the productivity of what they *thought* was the relevant metric (number of printouts) rather than what was *actually* necessary to improve productivity (immediacy of relevant and accurate information). Introducing technologies to tasks where there appear to be direct efficiency benefits, without considering if these tasks are superfluous, or at least *why* these tasks may benefit from efficiency improvements, is liable to lead to efficient inefficiency.

The issue of information handling and overload are likely to be particularly pronounced in the case of AI. Consider an illustrative example from a recent study by researchers from the software consultancy GitClear. They analysed several million lines of publicly available computer code and found that since 2022, while the *amount* of code written has increased significantly, so too has the amount of code 'churn' – editing and rewriting of code – leading to an overall decline in the apparent quality of code written (Harding and Kloster, 2024). GitClear attribute such effects to the rise of generative AI tools in programming. These tools can enable programmers to write more code, and thus they *appear* more productive. But much AI-generated code requires substantial editing from programmers, undermining any claims that AI is *actually* increasing productivity. In this instance, AI displaces productive programmers and increases the number of 'duct-tapers' (Graeber, 2018) who are charged with sorting through more inefficient outputs (i.e., broken code). Ironically, a human input remains important in making the code useable and an aid to productivity.

Bounded rationality must also be understood in the context of competing stakeholder pressures (Cyert and March 1963). On the one hand, shareholders may pressure senior managers to introduce AI technology quickly to reduce costs and signal that the organisation is at the cutting-edge (Acemoglu and Johnson, 2023). On the other hand, where they are present, unions may threaten senior managers with industrial action if they fear that the introduction of new technology will lead to job losses (Frey, 2020). The net result of these external influences is additional cognitive strain on senior managers when trying to identify superfluous tasks and the productive uses of AI technology.

As a result, some senior managers may settle for modest rather than maximum efficiency gains from adopted technology. For instance, effectively deploying AI may result in conflicts with junior managers and workers given automation fears (Brynjolfsson and McAfee, 2014), while the task of identifying effective use-cases may also reduce the time senior managers can spend pursuing their own organisational goals (Simon, 1947). Because of this, senior managers may be satisfied with some relative gains by deploying AI in 'easier' and less contested parts of the organisation's operations, rather than seeking out absolute

productive uses. Similarly, for shareholders who simply want to see AI deployment and (presumably) short run increases in organisational value given current popular interest, they may not commit the requisite time and effort to check whether AI is being used to carry out superfluous tasks (Acemoglu and Johnson, 2023; Veblen, 2016). Organisational conflicts and competing motives can create the conditions for satisficing behaviour, and thus efficient inefficiency (Simon, 1947).

A further example can help to illustrate the points discussed above. Some have suggested generative AI may be deployed to transcribe and summarise workplace meetings (Jaffe *et al.*, 2024; Levy, 2024). Similarly, Apple have recently announced plans to integrate 'smart reply' into their emailing service, allowing generative AI to automatically reply to emails (Zitron, 2024). These applications suggest efficiency gains may come through speeding up interpersonal and ideational interactions, freeing up employees to allocate their time more productively.

However, the productive benefits of a meeting typically come through the interpersonal exchange of ideas. The crafting of an email brings similar benefits – for example, it creates the opportunity to reflect on and understand the position of the respondent. If attendance at a meeting can be replaced with a summary of the meeting, the meeting itself may not be worthwhile. Likewise, if an email does not benefit from being personally written, the value of the email itself may be questioned. In each instance, AI would simply automate the *performance* of these workplace activities in a manner comparable to functional stupidity (Alvesson and Spicer, 2016; Alvesson and Stephens, 2024; Buchanan and Badham, 2020), suggesting that these activities – in some instances – are not especially beneficial to an organisation to begin with. Indeed, it suggests that fewer meetings and emails (and more human interaction) might actually add to overall productivity (also see Jaffe *et al.*, 2024).

The above applications of AI also fit into the information technology infrastructure that most organisations *already* have. This makes them *easy* to deploy, and thus enables senior managers to avoid the potential conflicts which may arise through a more disruptive (though potentially more productive) strategy for AI deployment. In this way, they confirm the above point about AI being used for easier rather than the most optimal applications.

Conflict between stakeholders can also create organisational uncertainty, which further frustrates the identification and elimination of superfluous tasks, and thus magnifies efficient inefficiency. A particular case to consider is the relationship between senior managers and junior managers (again, we focus on this case for illustrative purposes, in the knowledge that it does not reflect the management structure of all organisations). Junior managers necessarily know more about their team than senior managers and may be able to influence senior managers' understanding of the team through their control of the knowledge they possess (Simon, 1947). For example, junior managers may exaggerate the complexity of running their team in order to retain their control over it. Indeed, they may use exaggeration as a way to forestall the introduction of an automating technology. They may convince senior managers that it would be better to stick with them than seek more automated solutions that jeopardise their status and potentially their jobs. At the same time, they may create more superfluous tasks that make themselves appear busier and less dispensable (Faulconbridge *et al.* 2023; Lupu and Rokka, 2022).

To the extent that senior managers struggle to identify superfluous tasks and junior managers have incentives to create them, there will be a strong likelihood that AI technology will be deployed in ways which embed rather than reduce efficient inefficiency. The simple but pervasive problem of bounded rationality makes this outcome inevitable. To be sure, there is scope for senior managers to learn how to deploy AI technology more efficiently (Romer, 1990, 1989; Simon, 1991). Though, one should be cautious in stating that, in the long run, efficient inefficiency can be minimised or even eliminated. Rather, given the deep-seated nature of bounded rationality, efficient inefficiency is likely to endure, even in the context of rapid advances in AI.

4. Implications for practitioners

The below discussion draws together some strands of the preceding Sections, examining some implications for practitioners in terms of strategies to counteract efficient inefficiency. Bounded rationality cannot be overcome – it is an inevitable feature of the human condition and a fundamental feature of modern organisations (Simon, 1981). Instead, counters to efficient inefficiency must emerge from reducing the factors that frustrate managerial decision-making under bounded rationality (Simon, 1991, 1947). This is easier said than done. For instance, decisions around how to deploy technology will always be subject to some uncertainty, because they depend on future events with unpredictable and unknowable outcomes. This is particularly true of new technologies, such as emerging AI technology (Acemoglu and Johnson, 2023), where its scope and effects remain highly uncertain. Thus, technology deployment will typically be subject to some degree of satisficing, leaving open the possibility of efficient inefficiency.

Nevertheless, some of the difficulties in identifying superfluous tasks may be addressed through restructuring organisations to distribute the 'cognitive load' of managerial decisions. Furthermore, organisations could also be restructured to align stakeholder interests and utilise disparate knowledge through more democratic mechanisms. These strategies broadly align with holacracy perspectives on organisational structure, which emphasise flatter organisational arrangements, dispersed decision-making and autonomous working arrangements (e.g., Bernstein *et al.*, 2016; Laloux, 2014; Mintzberg, 1979; Robertson, 2015; Schell and Bischof, 2021). We also consider whether AI itself – as a tool for organising and manipulating information – could potentially support measures to tackle efficient inefficiency.

4.1. Distributed Decision-Making

Firstly, senior managers could embrace flatter organisational structures by delegating some decision-making power to junior managers as well as workers. The immediate benefit of this strategy would be to reduce the amount of information that senior managers have to deal with, and so the frequency with which they must rely on satisficing strategies (Lee and Edmondson, 2017). Bounded rationality theories of the firm can encourage one to imagine organisations as computers, with the organisation's employees representing its 'computational power' (Simon, 1981). Where senior managers retain decision-making power over technology deployment, they are under-utilising the computational power of their organisation.

A flatter organisational structure better utilises the *quantity* of decision-making power the organisation already has (Cyert and March 1963). It also promises to improve the *quality* of decision-making by increasing the diversity of information entering into organisational decisions (Simon, 1947). For instance, incorporating a worker's first-hand experience of a production problem, rather than a junior manager's second-hand experience, or a senior manager's third. The benefits of reducing the 'cognitive load' on senior managers, improving the quantity and quality of information used in decision-making and thus the quality of satisficing decisions about technology adoption, may offset any costs arising from junior managers and workers using their greater discretion to shirk.

One practical example to illustrate some of the points made above comes from Jeffrey Bezos (2016), the former CEO of Amazon, Inc., in his 2016 letter to shareholders. In it, he discusses a strategy he dubs 'disagree and commit'. It emphasises delegating decision-making authority down to more specialised junior managers who will typically be more knowledgeable about internal practices than senior managers. Bezos (2016) encourages such delegation even when there are disagreements between senior and junior managers, as in any disagreement he argues a more knowledgeable junior manager is more likely to be correct. We would add that in addition to 'disagree and commit' distributing the cognitive load of senior managers more evenly

throughout an organisation, it also challenges managerial motives for control, which, as argued above, often result in the production of superfluous tasks.

Such a simple strategy is not a panacea to some of the information processing problems discussed above, and distributed decision-making is not a prominent feature of Amazon when considering its many layers of management, not to mention its centralised control systems in its many distribution centres. Nevertheless, approaches like ‘disagree and commit’ recognise the role of distributed knowledge and information processing throughout an organisation, and the potential benefits of harnessing organisational capabilities through flatter organisational structures.

4.2. Aligning stakeholder interests

The above recommendation of distributed decision-making has reasonable challenges. For instance, junior managers may choose to use their delegated decision-making power to pursue their own interests, which may not be in the interests of senior managers, or the organisation as a whole. Restructuring organisations so that the interests of organisational stakeholders align – particularly senior and junior managers – is thus also important (Simon, 1991). The problem of agency highlighted above also carries further weight. Uncertainty about the actions of senior and junior managers may make it difficult to identify superfluous tasks. Junior managers’ uncertainty about the actions of senior managers may cause them to withhold useful information and to create superfluous tasks (Lupu and Rokka, 2022; Spicer, 2013). Senior managers’ uncertainty about the nature of a junior manager’s team may also mean that they struggle to determine superfluous from productive tasks (Egidi, 1997).

These sources of uncertainty may be countered by ensuring certain protections and benefits for junior managers, such as job security and the scope for progression. Additionally, senior managers could ensure the potential benefits of AI, such as reduced working hours, are shared with junior managers (and workers), allowing the latter to retain or improve the quality of their work, while doing less of it. Sociological perspectives that emphasise the motivating role of status would likely support these recommendations (Graeber, 2018, 2015), as would behavioural accounts (Cyert and March 1963), and some managerial theories that emphasise the importance of aligning junior manager interests with those of senior managers to ensure outcomes preferable to the organisation (Friedman, 1987).

It is interesting to consider that ‘aligning stakeholder interests’ extends to other stakeholders, such as workers, and could lead to quite radical organisational restructuring. Just as there exists conflict between senior and junior managers, so there exists conflict between junior managers and workers. Workers who fear redundancy because of automating technologies like AI have little incentive to disclose whether their work is superfluous or not and face similar incentives to invent superfluous tasks (Spicer, 2013). Like senior managers, junior managers may also struggle to identify these tasks given their degree of separation from them (Laloux, 2014; Robertson, 2015). Job security, the sharing of economic gains, and further deployment of ‘disagree and commit’ style strategies at this lower level, may help to align the interests of junior managers and workers and create scope for the more efficient use of technology.

Furthermore, all the stakeholders discussed thus far – senior managers, junior managers, and workers – may have objectives that conflict with those of shareholders. Therefore, one solution may be to give these stakeholders a direct ownership stake in the organisation. A radical solution would be to remove shareholders and make all employees joint owners. Some research into flatter, more collectively managed organisations (including cooperatives) suggest these organisations are highly adaptable in periods of technological change (e.g., Ackermann et al., 2021).

For example, Laloux (2014) discusses technology adoption decisions

at FAVI, a gear-box manufacturing company. At FAVI, teams autonomously organise to meet orders for different clients and set manufacturing plans based on collective agreements, while team members share risks and rewards. FAVI’s organisational structure thus incentivises members to find productive uses for technology, while *disincentivising* the creation of superfluous tasks in the first instance. New technologies are frequently presented to these autonomous teams by dedicated technology scouts but are only adopted when teams collectively determine the technology will enhance their processes. As Laloux (2014) reports, FAVI’s innovation-led approach has resulted in it establishing a market-leading position. Other studies have found worker-owned firms are more productive, and more likely to benefit from greater productivity gains from technological adoption, than conventional capitalist firms (Craig and Pencavel, 1995; Fakfakh et al., 2012; Pérotin, 2016), confirming the economic advantage of forms of shared or collective management.

Aligning stakeholder interests does not mean that superfluous tasks would cease to materialise, or that stakeholder groups would always work to reduce informational limits and uncertainty. A worker who especially likes the tasks they perform may fear automation of those tasks, even if they have a job guarantee and a share of the automation benefits (Laloux, 2014). Managers, both senior and junior, may still create superfluous tasks for the purposes of control and surveillance (Beer, 1993; Tiratelli, 2022). Inevitable conflicts coupled with fundamental uncertainty will still require managers to engage in satisficing strategies to manage their bounded cognitive resources. But, as argued above, reducing incentives to create informational uncertainty through fostering shared interests may improve an organisation’s abilities to identify superfluous tasks, reducing the risk of efficient inefficiency.

4.3. Democratic Decision-Making

Another complementary yet radical organisational strategy, given typical decision-making structures in capitalist firms, is to introduce more democratic decision-making into organisations. The major advantage of democratic decision-making as a means of enhancing cooperation within organisations is that it allows disparate knowledge to be efficiently synthesised into a collective decision (Simon, 1981, 1947), thus emulating the advantages of market mechanisms (Hayek, 1945).

Organisations often contain significant tacit knowledge (knowledge which cannot be easily transformed into data), local knowledge (knowledge with high specificity to an organisation), and specialist knowledge (scarce knowledge of a generic process). All can be captured in what transaction costs economics calls ‘human asset specificity’, or the value of a specific employee’s knowledge to an organisation (Williamson, 1999). Through democratic engagement of all employees (senior and junior managers, workers, etc.) in decisions surrounding technology adoption and deployment, organisations may efficiently bring together these valuable knowledge resources (Farrell and Shalizi, 2012). In instances of high uncertainty, such as decisions concerning the adoption of emergent technology like AI, democratic approaches which efficiently integrate many different perspectives may be particularly valuable (Muthukrishna, 2023).

Collective, democratic decision-making may also serve to avoid the duplication of expertise. For instance, AI technology is often sold as ‘a service’ (Cobbe and Singh, 2021; Lins et al., 2021), with external consultants recruited to supplement knowledge ‘gaps’ (e.g., Amazon Web Services, 2023; SAP, 2023). Yet, such ‘gaps’ might only exist because senior managers are ignorant of existing organisational knowledge (Lee and Edmondson, 2017). If so, outsourced expertise would be a duplication of knowledge and represent a source of superfluous work. Therefore, regardless of the decisions taken, democratic decision-making may in itself be an efficient means of ‘auditing’ the knowledge an organisation has, avoiding duplication and lowering the risk of efficient inefficiency.

Democratic processes, such as worker representation at board level and works councils, may have added economic benefits, including promoting knowledge diffusion about the use and application of new technology like AI throughout an organisation (Ackermann et al., 2021). Research has found that organisations with more dispersed knowledge of vital processes are more innovative and adaptive to market conditions in the long run (Rodan and Galunic, 2004). Insofar as democratic decision-making can encourage the sharing and diffusion of knowledge, its proliferation may help to enhance innovation and adaptability (Farrell and Shalizi, 2012). Its benefits may be magnified, in turn, by moves to share ownership in the ways described above.

4.4. AI applications

While the above recommendations emphasise restructuring organisations in order to reduce efficient inefficiency when introducing new technologies, it is worth noting that technologies themselves can have transformative effects on organisational structures (Acemoglu and Johnson, 2023). As noted above, the introduction of various automating technologies into industries has historically prompted the reorganisation of groups within organisations (Frey, 2020), and worries about the scope and extent of job losses linked to automation continue to shape decision-making around technology to this day (Brynjolfsson and McAfee, 2014).

AI technology, in particular, has the potential to catalyse – and support – the transformation of organisations. Medina (2014, p. 29), writing on the history of cybernetic management, notes that “computer technology could help organise the parts of [a firm] into a better-functioning whole... Computers [do] not need to reinforce existing management hierarchies and procedures; instead, they could bring about structural transformation in a company.” This is to say, AI technology can be used to understand how parts of the organisation itself function, interact, and contribute to its overall objectives, which in turn could support senior managers in reorganising work and management to achieve superior outcomes. Emerging proposals explore some aspects of this idea.

Firstly, AI technology has long been proposed as a tool for supporting boundedly rational decision-makers to make better decisions. As earlier as 1987, Simon (1987a, 1987b) speculated on the potential collaboration between expert decision-makers and AI systems. He suggested that decision-makers in organisations often make worse decisions because the structure of the organisations means they must assimilate too much information – a suggestion we have already covered above. Simon (1987a, 1987b) suggested training an AI system to learn precisely what information is actually relevant to successful decision-making, allowing organisations to filter out useless or redundant information. This idea is likely an evolution of the problem, and solution, identified by Simon (1981) in his discussion of the diplomatic office.

Similarly, emerging work suggests that AI can be used to address human decision-making processes and identify biases in these processes that may lead to suboptimal outcomes. Much of this work focuses on decision-makers in the public domain, such as judges (Kleinberg et al., 2018), but some has also examined decisions with more commercial relevance, such as decisions by doctors to refer patients to potentially expensive, but likely unnecessary, medical tests (Mullainathan and Obermeyer, 2022). These applications may allow AI technology to transform organisations, potentially fostering more ‘rational’ decision-making that could yield economic benefits. Nevertheless, as argued above, superfluous tasks and thus efficient inefficiency are likely to arise, even with better decision-making. This is because these tasks are also linked to organisational conflicts between members of an organisation, which require non-technological responses and solutions.

Secondly, AI technology could be utilised as a tool for decentralising and redistributing knowledge, and as a means for encouraging more participatory – if not democratic – decision-making within organisations. Consider, for instance, large language models. These provide easy

access to knowledge and enable employees at different levels of the organisation to participate in decision-making. Specian (2023) has suggested that these models can make expertise more accessible to non-experts, offering an alternative means of accessing expertise and thus supporting new, democratic organisational mechanisms. Others (Bouschery et al., 2023; Broekhuizen et al., 2023) have emphasised how AI technology may be used to support organisational innovation through bringing together disparate organisational knowledge, to support decision-makers, while Ermakova and Frolova (2021) have argued that AI technology can be used to resolve conflicts within organisations. All such applications point to the potential of AI technology to transform the channels through which members of organisations communicate, learn, and take decisions, which may support a programme of organisational restructuring to overcome efficient inefficiency.

Yet, some caution must be shown here. As Specian (2024, 2023) acknowledges, large language models could also be used within organisations to centralise knowledge, giving senior managers greater control over information, potentially to the detriment of other members of the organisation (though, not necessarily to managerial motives; see Marglin, 1974 for a classic study on how managers can use technology for the purposes of control). Thus, once more, while AI technology has the scope to facilitate distributed decision-making and democratic processes, it offers no guarantee that these outcomes will be achieved. Rather, the use of AI and its effects (including on efficiency) very much depend on who controls and directs it.

4.5. Summary

The complex nature of managerial decision-making under conditions of bounded rationality mean it is unlikely that managers (both senior and junior) can identify and eliminate all superfluous tasks within an organisation. The risk of technology (including AI) being deployed to superfluous, rather than productive, tasks is likely to remain. This will reduce the economic benefits that organisations can achieve even with the latest technology.

Nevertheless, understood as a problem of bounded rationality, restructuring organisations to reduce the cognitive and uncertainty challenges that make it harder to identify superfluous tasks should help to reduce the extent and magnitude of efficient inefficiency, as should restructuring which disincentivises the creation of such tasks in the first instance. Inevitably, the benefits of restructuring in terms of combatting efficient inefficiency must be weighed against the costs of restructuring (in time and money), which may be substantial. In some instances, one might also speculate that technology such as AI could support restructuring efforts directly, though this will depend on who directs it and how it is used.

We would add a final, crucial point. The notion of restructuring does not mean demanding that all stakeholders hold common interests in the goal of profit maximisation. It is not a matter of ensuring that senior managers get junior managers and workers to act in efficient ways so that profit is maximised (Laloux, 2014). Rather, we would argue that restructuring is warranted precisely because members of organisations will always have a diverse set of goals, and that inefficiencies arise when these goals are ignored. Ensuring organisations represent and meet a plurality of aims is important in enhancing efficiency (Simon, 1947). In addition, pushing for restructuring can mean thinking beyond profit maximisation as the sole or primary goal of organisations. Thinking through how efficient inefficiency can be overcome, in short, can entail reimagining how organisations are directed. This extends to how AI is directed and how its economic rewards are distributed.

5. Implications for theory

To the extent that efficient inefficiency encompasses consideration of technology adoption, it is worthwhile briefly examining some of the predictions which arise from it. One prediction has already been

discussed, namely that efficient inefficiency resulting from bounded rationality implies organisations adopting more democratic structures – ones that facilitate inclusive and collective decision-making – will likely be more productive. As mentioned above, some studies of cooperative organisations (Craig and Pencavel, 1995; Fakfakh et al., 2012; Pérotin, 2016) support this prediction.

Efficient inefficiency also predicts that the productivity benefits of introducing new technology should be less substantial in the short run than would otherwise be expected. Within the economics of technology literature, this prediction is known as the Solow paradox (Solow, 1987) or the productivity paradox (Brynjolfsson, 1993). Broadly, the paradox describes the often-observed lag (sometimes several decades) between technology being introduced and productivity increasing (Frey, 2020). Efficient inefficiency can be applied here to understand how this paradox might persist, even in the context of more rapid developments in AI.

One explanation of the Solow paradox concerns so-called ‘human capital’. The argument is that new technology can only be used productively if workers develop new skills. For instance, early computers were generally useless without someone who knew how to interface with them (e.g., coding computer programmes). It takes time – perhaps a generation – for requisite skills to be common enough within the workforce for the benefits of new technology to be realised. The lag between the initial introduction of a technology and the effect of the technology on firm productivity may be attributed to the time taken for workers (as well as managers) to acquire the skills to productively use it (Frey, 2020; Romer, 1990, 1989).

We do not disagree with the human capital account of the Solow paradox. However, human capital is distinct from efficient inefficiency as an explanation of it. The human capital perspective focuses on technology, regardless of the task to which technology is applied. It assumes implicitly that before knowledge of a technology is widely held, managers can accurately determine how it should be used. This is understandable insofar as some economic perspectives, as mentioned above, are sceptical about the competitive viability of inefficient firms and thus the longevity of superfluous tasks (Alchian, 1950; Friedman, 1953; Simon, 1981). However, by focusing on the task, rather than on the technology, efficient inefficiency makes no assumption that managers have prior knowledge of how to use technology productively, nor does it deny the possibility that organisations feature superfluous tasks. As such, efficient inefficiency has theoretical benefits in elucidating the Solow paradox, either as a complement to the human capital explanation, or as an alternative account in its own right.

Expertise in the use of a technology, such as AI, is different to expertise in specific organisational processes. A senior manager who identifies an opportunity for efficiency improvements may lack the human capital to exploit this opportunity, perhaps because they are ignorant of the availability and applicability of relevant technologies. But equally, a senior manager with knowledge of these technologies may be wholly ignorant of worthwhile use cases, meaning that this knowledge fails to translate into higher efficiency. In practice, it may take a long time for managers to acquire both the human capital to know how to use a technology and the organisational knowledge to identify how a technology ought to be used to avoid efficient inefficiency.

Efficient inefficiency is therefore a useful contribution to the study of the economics of technology, as much as to managerial decision-making and organisational theory. In particular, it offers a new perspective on the Solow paradox. If the AI revolution currently underway follows the pattern of past technologies in failing to have an immediate productivity benefit, efficient inefficiency will be an important candidate in explaining this outcome.

6. Future perspectives on efficient inefficiency for practitioners

Future efforts to investigate efficient inefficiency will benefit from empirical enquiry into technology adoption decisions within real world

organisations. This will include case-studies of how AI has been adopted and its effects on tasks. Interviews with managers (both senior and junior) might also be useful in gauging their views on the efficiency of existing tasks and their responses to new technology as a potential and actual mechanism for raising efficiency.

There are obvious challenges here. The very notion of efficient inefficiency rests on data that is not easily obtainable. Much of the literature examining superfluous work and related phenomena struggle to gather data and are frequently criticised for their speculative nature (e.g., Tiratelli, 2022). Organisations that can effectively measure, and thus identify, superfluous tasks are likely to eliminate such tasks, eroding one source of empirical data. Studies of superfluous work come to rely on anonymous self-reported data from members of organisations (e.g., Graeber, 2018), or consultants, whose external view allows them to more critically comment on internal processes (e.g., Medina, 2014; Simon, 1981). These data and the views of consultants, however, can be prone to error and lack objectivity. Finally, as mentioned above, organisations may not even realise they are suffering from efficient inefficiency. If a superfluous task is considered productive, the relative benefit of a technology deployed to it may be construed as a worthwhile, absolute benefit (see, for instance, Jaffe et al., 2024). It may be held up as an exemplar of effective technology deployment, rather than as an example of efficient inefficiency. This raises challenges for empirical work aimed at studying its nature and extent as well as ways of managing it.

One approach may be to view efficient inefficiency taxonomically, distinguishing between different types of inefficiency, and thus between different organisational contexts in which it arises. For instance, the primary definition of efficient inefficiency presented above suggests that it occurs when superfluous tasks are supplemented with new technology. However, we have also presented examples where technology supplements productive tasks to the point of undermining productivity (e.g., teleprinters in the diplomatic office), where technology displaces productive elements, potentially turning a productive task into a superfluous one (e.g., ‘smart reply’ technology), and where technology creates ‘duct-taping’ work, undermining any potential new economic benefits it may yield (e.g., coding co-pilot tools). For practitioners and researchers alike, investigating specific types of efficient inefficiency and their organisational origins may be more advantageous than treating the concept necessarily as a singular phenomenon.

Further investigation of the mechanisms through which efficient inefficiency can arise will also be crucial for developing practical recommendations. This paper has suggested efficient inefficiency is a product of bounded rationality within organisations, and that many instances of efficient inefficiency are likely to arise when organisational problems are mistaken as technical ones. A potentially useful future direction for practice as well as research is therefore to understand how and why these mistakes occur. While a useful framework for considering individual and organisational behaviour, bounded rationality can suffer from being overly general, which in turn undermines efforts to develop strategies for changing managerial behaviour and improving organisational performance (Simon, 1981, 1947). Moving from theory to practice may then entail looking beyond bounded rationality as a concept and towards the different ways in which managers actually make decisions, including but not limited to considering the wider sociological and economic market pressures which shape technology adoption decisions (e.g., Acemoglu and Johnson, 2023).

At a more direct level, there is work to be done in ensuring that technology like AI can be used to improve efficiency while enhancing well-being. Efficiency should not be at the expense of well-being and the challenge will be in building organisations and an economy that can channel new technology in ways that benefit people, not just profit. From the arguments made above, this will entail some radical rethinking of organisations, including shifts in ownership. Indeed, without considerations of how the benefits of technology may be distributed throughout organisations and society at-large, those who stand the least

to benefit from it may also be disposed to undermine the technology's effective deployment.

7. Conclusion

We have argued that technology is efficiency-enhancing when used to perform productive tasks more efficiently. However, if deployed to perform unproductive, or *superfluous*, tasks, the economic benefits of technology are reduced, leading to what we call efficient inefficiency. Efficient inefficiency has important implications for organisational theory and managerial decision-making and explains why new technology can fail to substantially enhance efficiency, at least in the short run. Given recent excitement around AI, we have focused on this technology specifically, and considered how efficient inefficiency may apply to its deployment in organisations.

Drawing on sociological, cybernetic, and behavioural perspectives, we have argued that superfluous tasks exist in most (if not all) organisations. These tasks arise for various reasons, such as managerial motives for status and control, or as defensive strategies specifically deployed by junior managers. Deploying technology in productive ways depends on managers effectively identifying superfluous tasks. However, they may struggle to identify these tasks because of bounded rationality. Dividing management into senior managers, who take decisions, and junior managers, who respond to decisions, we have argued that senior managers often lack the informational resources to determine superfluous from productive tasks, and that junior managers can create substantial uncertainty about tasks, further hindering effective decision-making. Where superfluous tasks exist and cannot be readily identified, there is a clear risk of efficient inefficiency arising and persisting.

As a response, we suggest organisations could be restructured to reduce the factors which make identifying superfluous tasks difficult. Flatter organisational structures and shared decision-making about technology deployment may reduce the informational burden that falls on senior managers. Aligning the organisational interests of stakeholders may also encourage the disclosure of superfluous tasks and discourage their creation in the future. Democratic decision-making may efficiently utilise the local, specialist, and tacit knowledge contained within organisations, allowing senior managers to make more informed decisions about how technology should be used, while giving more junior managers and workers greater autonomy over their work. In addition to these recommendations, we have considered how AI as an information management technology may support the restructuring of organisations to overcome bounded rationality effects and thus efficient inefficiency.

In terms of implications for theory, we have argued that efficient inefficiency contributes to an explanation of the often-observed lag between technological developments and efficiency improvements. In this way, we have argued efficient inefficiency furthers understanding of the economics of technology as well as of the origins of economic stagnation in a world of seemingly rapid technological progress. We have also identified some future implications for practitioners to consider, and some future avenues for investigating efficient inefficiency.

We urge scholars and practitioners concerned with the use of new technology in organisations to view deployment decisions not as technical matters, but as matters with high social and political import. The central question is rarely whether a technology can perform a task, but whether the task is worth performing at all. Those who ignore this question risk overlooking the problem of efficient inefficiency. They also miss the opportunity to rethink the problem of organisation and to find novel solutions to it.

CRedit authorship contribution statement

Stuart Mills: Writing – review & editing, Writing – original draft, Conceptualization. **David A. Spencer:** Writing – review & editing, Writing – original draft, Conceptualization.

References

- Acemoglu, D. (2024). *The Simple Macroeconomics of AI*. MIT Shaping the Future of Work Institute. [Online] [Date accessed: 18/06/2024]: https://shapingwork.mit.edu/wp-content/uploads/2024/04/Acemoglu_Macroeconomics-of-AI_April-2024.pdf.
- Acemoglu, D., & Johnson, S. (2023). *Power and progress: our thousand-year struggle over technology and prosperity*. UK: Basic Books.
- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: how technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30.
- Ackermann, M., Schell, S., & Kopp, S. (2021). How Mercedes-Benz addresses digital transformation using Holacracy. *Journal of Organizational Change Management*, 34(7), 1285–1299.
- Ahuja, A. (2023). *'Generative AI is sowing the seeds of doubt in serious science'* The Financial Times. [Online] [Date accessed: 07/07/2023]: <https://www.ft.com/content/e34c24f6-1159-4b88-8d92-a4bda685a73c>.
- Alchian, A. A. (1950). Uncertainty, evolution, and economic theory. *Journal of Political Economy*, 58(3), 211–221.
- Alvesson, M., & Stephens, A. (2024). "Is it worth doing this or is it better to commit suicide?": On ethical clearance at a university'. *Human Relations*, 1–30.
- Alvesson, M., Spicer, A. (2016) *The Stupidity Paradox: The Power and Pitfalls of Functional Stupidity at Work* Profile: UK.
- Alvesson, M., & Spicer, A. (2012). A stupidity-based theory of organizations. *Journal of Management Studies*, 49(7), 1194–1220.
- Amazon Web Services (2023) *'Machine Learning on AWS'* Amazon Web Services. [Online] [Date accessed: 16/07/2023]: <https://aws.amazon.com/machine-learning/>.
- Audreusch, D. B., & Belitski, M. (2023). Digitization, resource mobilization and firm growth in emerging industries. *British Journal of Management*. <https://doi.org/10.1111/1467-8551.12769>
- Autor, D. H. (2015). Why are there still so many jobs? The History and Future of Workplace Automation. *Journal of Economic Perspectives*, 29(3), 3–30.
- Beer, S. (1993) *'Designing Freedom'* Anansi: Canada.
- Beer, S. (1979) *'Decision and Control'* Wiley: UK.
- Bernstein, E., Bunch, J., Canner, N., & Lee, M. (2016). Beyond the holacracy hype. *Harvard Business Review*, 94(7–8), 38–39.
- Bezos, J. (2016) *'2016 Letter to Shareholders'* Amazon, Inc. [Online] [Date accessed: 16/11/2023]: https://s2.q4cdn.com/299287126/files/doc_financials/annual/2016-Letter-to-Shareholders.pdf.
- Bouschery, S. G., Blazeovic, V., & Piller, F. T. (2023). Augmenting human innovation teams with artificial intelligence: Exploring transformer-based language models. *Journal of Product Innovation Management*, 40, 139–153.
- Broekhuizen, T., Dekker, H., de Faria, P., Firk, S., Nguyen, D. K., & Sofka, W. (2023). AI for managing open innovation: Opportunities, challenges, and a research agenda. *Journal of Business Research*, 167, e114196.
- Brynjolfsson, E. (1993). The productivity paradox of information technology. *Communications of the ACM*, 36(12), 66–77.
- Brynjolfsson, E., McAfee, A. (2014) *'The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies'* W. W. Norton: USA.
- Buchanan, D., & Badham, R. (2020). *Power, politics, and organizational change*. UK: Sage.
- Cambon, A., Hecht, B., Edelman, B., Ngwe, D., Jaffe, S., Heger, A., Vorvoreanu, M., Peng, S., Hofman, J., Farach, A., Bermejo-Cano, M., Knudsen, E., Bono, J., Sanghavi, H., Spatharioti, S., Rothschild, D., Goldstein, D. G., Kalliamvakou, E., Cihon, P., Demirer, M., Schwarz, M., Teevan, J. (2023) *'Early LLM-based Tools for Enterprise Information Workers Likely Provide Meaningful Boosts to Productivity'* Microsoft. [Online] [Date accessed: 13/08/2024]: <https://www.microsoft.com/en-us/research/ch/uploads/prod/2023/12/AI-and-Productivity-Report-First-Edition.pdf>.
- Cobbe, J., & Singh, J. (2021). Artificial intelligence as a service: legal responsibilities, liabilities, and policy challenges. *Computer Law and Security Review*, 42, e105573.
- Craig, B., & Pencavel, J. (1995). Participation and productivity: a comparison of worker cooperatives and conventional firms in the plywood industry. *Brookings Papers on Economic Activity: Microeconomics*, 1995(1), 121–174.
- Cyert, R. M., & March, J. G. (1963). *A behavioral theory of the firm*. USA: Prentice-Hall.
- Dean, E., Dadzie, R. B., & Pham, X. (2022). The instinct of workmanship and the incidence of bullshit jobs. *Journal of Economic Issues*, 56(3), 673–698.
- Delfanti, A., & Frey, B. (2021). Humanly extended automation or the future of work seen through amazon patents. *Science, Technology and Human Values*, 46(3), 655–682.
- Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., Lakhani, K. R. (2023) *'Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality'* SSRN. [Online] [Date accessed: 07/01/2024]: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321.
- Dodgson, M. (1991). Technology learning, technology strategy and competitive pressures. *British Journal of Management*, 2, 133–149.
- Dosi, G., & Marengo, L. (2007). On the evolutionary and behavioral theories of organizations: a tentative roadmap. *Organization Science*, 18(3), 337–345.
- Egidi, M. (1997). Reply to simon. In H. A. Simon (Ed.), *An Empirically Based Microeconomics*. UK: Cambridge University Press.
- Ermakova, E. P., & Frolova, E. E. (2021). Using artificial intelligence in dispute resolution. In A. O. Inshakova, & E. E. Frolova (Eds.), *Smart Technologies for the Digitisation of Industry: Entrepreneurial Environment*. USA: Springer.
- Fakfakh, F., Pérotin, V., & Gago, M. (2012). Productivity, capital and labor in labor-managed and conventional firms. *Industrial and Labor Relations Review*, 65(4), 847–878.
- Faraj, S., & Pachidi, S. (2021). Beyond Uberization: The co-constitution of technology and organizing. *Organization Theory*, 2(1). <https://doi.org/10.1177/2631787721995205>

- Farrell, H., Shalizi, C. (2012) 'Cognitive Democracy' [Online] [Date accessed: 13/08/2024]: https://crookedtimber.org/wp-content/uploads/2012/05/cognitive_democracy_may2012.pdf.
- Faulconbridge, J., Sarwar, A., & Spring, M. (2023). How professionals adapt to artificial intelligence: the role of intertwined boundary work. *Journal of Management Studies*. <https://doi.org/10.1111/joms.12936>
- Frey, C. B. (2020). *The technology trap: capital, labor, and power in the age of automation*. UK: Princeton University Press.
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: how susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280.
- Friedman, A. L. (1987). *Industry and labour: class struggle at work and monopoly capitalism*. UK: Macmillan.
- Friedman, M. (1953). *Essays in positive economics*. USA: University of Chicago Press.
- Graeber, D. (2018). *Bullshit jobs: a theory*. UK: Penguin Books.
- Graeber, D. (2015). *The utopia of rules: on technology, stupidity, and the secret joys of bureaucracy*. UK: Melville House.
- Harding, W., Kloster, M. (2024) 'Coding on Copilot: 2023 Data Shows Downward Pressure on Code Quality' GitClear. [Online] [Date accessed: 17/06/2024]: https://www.gitclear.com/coding_on_copilot_data_shows_ais_downward_pressure_on_code_quality.
- Hayek, F. A. (1945). The use of knowledge in society. *The American Economic Review*, 35(4), 519–530.
- Hodgkinson, G. P., Burkhard, B., Foss, N. J., Grichnik, D., Sarala, R. M., Tang, Y., & van Essen, M. (2023). The heuristics and biases of top managers: past, present, and future. *Journal of Management Studies*, 60(5), 1033–1063.
- Karube, M., Numagami, T., & Kato, T. (2009). Exploring organisational deterioration: 'organisational deadweight' as a cause of malfunction of strategic initiatives in Japanese firms. *Long Range Planning*, 42(4), 518–544.
- Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., & Mullainathan, S. (2018). Human decisions and machine predictions. *The Quarterly Journal of Economics*, 133(1), 237–293.
- Laloux, F. (2014). *Reinventing organizations*. USA: Diateino.
- Lee, M. Y., & Edmondson, A. C. (2017). Self-managing organizations: exploring the limits of less-hierarchical organizing. *Research in Organizational Behavior*, 37, 35–58.
- Levy, S. (2024) 'A Wave of AI Tools Is Set to Transform Work Meetings' Wired Magazine. [Online] [Date accessed: 18/06/2024]: <https://www.wired.com/story/takin-g-baby-steps-toward-the-ai-meeting-singularity/>.
- Lins, S., Pandl, K. D., Teigeler, H., Thiebess, S., Bayer, C., & Sunyaev, A. (2021). Artificial intelligence as a service. *Business Information Systems Engineering*, 63(4), 411–456.
- Lupu, I., & Rokka, J. (2022). "Feeling in Control": optimal busyness and the temporality of organizational controls. *Organization Science*, 33(4), 1396–1422.
- Marglin, S. A. (1974). What do bosses do? The origins and functions of hierarchy in capitalist production. *Review of Radical Political Economics*, 6(2), 60–112.
- McGee, P., Murgia, M. (2023) 'Microsoft to add AI co-pilot to its Office software suite' The Financial Times. [Online] [Date accessed: 07/07/2023]: <https://www.ft.com/content/c74c7e48-d439-40d1-a42a-2cd9f1560edb>.
- Medina, E. (2014). *Cybernetic revolutionaries: technology and politics in allende's chile*. USA: MIT Press.
- Mintzberg, H. (1979). The structuring of organizations. In D. Asch, & C. Bowman (Eds.), *Readings in Strategic Management*. USA: Springer.
- Mullainathan, S., & Obermeyer, Z. (2022). Diagnosing physician error: a machine learning approach to low-value health care. *The Quarterly Journal of Economics*, 137(2), 679–727.
- Muthukrishna, M. (2023). *A theory of everyone: who we are, how we got here, and where we're going*. UK: Basic Books.
- Perlow, L. A. (1998). Boundary control: the social ordering of work and family time in a high-tech corporation. *Administrative Science Quarterly*, 43(2), 328–357.
- Pérotin, V. (2016) 'What do we really know about worker co-operatives?' Co-operatives UK. [Online] [Date accessed: 16/07/2023]: https://www.uk.coop/sites/default/files/2020-10/worker_co-op_report.pdf.
- Robertson, B. J. (2015). *Holacracy: the revolutionary management system that abolishes hierarchy*. UK: Penguin Books.
- Rodan, S., & Galunic, C. (2004). More than network structure: How knowledge heterogeneity influences managerial performance and innovativeness. *Strategic Management Journal*, 25(6), 541–562.
- Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy*, 98(5), 71–102.
- Romer, P. M. (1989). Human capital and growth: Theory and evidence. *Carnegie-Rochester Conference Series on Public Policy*, 32, 251–286.
- SAP (2023) 'Generative AI Roadmap' SAP. [Online] [Date accessed: 07/07/2023]: <https://www.sap.com/products/artificial-intelligence/generative-ai.html>.
- Schell, S., & Bischof, N. (2021). Change the way of working. Ways into self-organization with the use of Holacracy: An empirical investigation. *European Management Review*, 19(1), 123–137.
- Schoemaker, P. J. H. (1993). Strategic decisions in organizations: rational and behavioural views. *Journal of Management Studies*, 30(1), 107–129.
- Simon, H. A. (1997). *An empirically based microeconomics*. UK: Cambridge University Press.
- Simon, H. A. (1991). Bounded rationality and organizational learning. *Organization Science*, 2(1), 125–134.
- Simon, H. A. (1987a). Two heads are better than one: the collaboration between AI and OR. *Interfaces*, 17(4), 8–15.
- Simon, H. A. (1987b). Making management decisions: the role of intuition and emotion. *The Academy of Management Executive*, 1(1), 57–64.
- Simon, H. A. (1981). *The sciences of the artificial* (2nd ed). USA: MIT Press.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118.
- Simon, H. A. (1947). *Administrative behavior*. USA: Macmillan.
- Solow, R. (1987) 'We'd better watch out' New York Times Book Review. [Online] [Date accessed: 19/07/2023]: <http://www.standardeconomist.com/pdf/misc/solow-computer-productivity.pdf>.
- Specian, P. (2024) 'Machine Advisors: Integrating Large Language Models into Democratic Assemblies' SSRN. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4682958.
- Specian, P. (2023) 'Large Language Models for Democracy: Limits and Possibilities' Technology and Sustainable Development 2023 Conference Paper. Available at: https://techandsd.com/files/specian_2023.pdf.
- Spencer, D. A. (2022) 'Making Light Work: An End to Toil in the Twenty-First Century' Polity: UK.
- Spicer, A. (2013). Shooting the shit: The role of bullshit in organisations. *M@nagement*, 5(16), 653–666.
- Stubbart, C. I. (1989). Managerial cognition: a missing link in strategic management research. *Journal of Management Studies*, 26(4), 325–347.
- Tiratelli, M. (2022). 'Where did the bullshit jobs go?'. *Catalyst A Journal of Theory and Strategy*, 6(3), 62–98.
- Veblen, T. (2016). *The theory of business enterprise*. UK: Scribner.
- Weber, M. (2019). *Economy and society*. USA: Harvard University Press.
- Williamson, O. (1999). Strategy research: Governance and competence perspectives. *Strategic Management Journal*, 20, 1087–1108.
- Williamson, O. (1981). The economics of organization: The transaction cost approach. *American Journal of Sociology*, 87(3), 548–577.
- Zitron, E. (2024) 'Let Tim Cook' [Online] [Date accessed: 18/06/2024]: <https://www.wheresyoured.at/untitled/>.

Stuart Mills is assistant professor of economics at the University of Leeds. His research focuses on decision-making and technology; on how technology shapes human behaviour; and on organisational decision-making concerning technology adoption and development.

David A. Spencer is Professor of Economics at the University of Leeds. His research lies in labour economics and political economy. He uses an interdisciplinary approach to engage with ideas from economics and other disciplines to explore and understand the nature of work.