



**Flinders
University**




Jeff Bleich Centre
for Democracy and
Disruptive Technologies



BEYOND CREATIVE DESTRUCTION

The Productive-Security State and the Productive-Social State: Competing Uses of the AI Dividend

Ottorino Cappelli



SEPTEMBER
2026



© Flinders University, 2026

Except where noted, this work is licensed under CC BY 4.0

Images sourced under licence from Canva are excluded from the CC BY licence.

Suggested Citation: Cappelli, O. (2026). "BEYOND CREATIVE DESTRUCTION: The Productive-Security State and the Productive-Social State: Competing Uses of the AI Dividend", Jeff Bleich Centre for Democracy and Disruptive Technologies, Flinders University, South Australia
DOI: <https://doi.org/10.25957/f7ye-r893>

jbc@flinders.edu.au

BEYOND CREATIVE DESTRUCTION

The Productive-Security State and the Productive-Social State: Competing Uses of the AI Dividend

Ottorino Cappelli¹

1. A Choice, Not a Destiny

In public debate on artificial intelligence, a sequence has taken hold that seems to admit no deviation. A machine learns to perform a task. The firm then discovers that it can produce with fewer people, who must look elsewhere for work while the state tries, if it can, to compensate the losers. This is the contemporary version of Schumpeter's creative destruction, accelerated until innovation and redundancy occur at the same moment. To its advocates, that is the unavoidable price of progress. A substantial part of the critics accepts the same premise and draws the opposite conclusion: if this is the progress promised by AI, it should be constrained and kept out of the workplace.

But the decisive step comes earlier. Eliminating a task and expelling the person who performs it are not the same event. What joins them is an organisational and political choice, supported by economic incentives and power relations that make it appear to be a mere effect of technology. Labour and innovation economists, together with scholars of organisation, have long distinguished between automation and augmentation. In the first case, technology replaces human labour in tasks previously entrusted to people. In the second, it strengthens their capabilities and opens up new activities. Daron Acemoglu and Pascual Restrepo have shown that automation produces a displacement effect, whereas the creation of new tasks in which labour retains an advantage produces a reinstatement effect. Erik Brynjolfsson has called the obsession with machines capable of imitating humans the Turing Trap: if the aim is to replicate the worker, the most likely outcome is replacement; if we build tools that extend human capabilities, the value of human experience also rises. (Acemoglu and Restrepo, 2019; Brynjolfsson, 2022).

The first empirical tests available confirm that this second path exists. In a large customer-support firm, a generative system that suggested responses while leaving operators free to modify or ignore them increased the number of issues resolved per hour by 15 per cent. The gain reached 30 per cent among less experienced workers, who used the tool to learn more quickly from the practices of the best-performing colleagues. It was not an artificial worker placed in their stead, but a device that spread through the organisation knowledge previously concentrated in a few people. The study does not warrant broad predictions and cautions that, in the long run, the same firm might use those gains to reduce headcount. It

¹ Ottorino Cappelli is a Research Fellow at the Jeff Bleich Centre for Democracy and Disruptive Technologies. He teaches Comparative Politics at the University of Naples L'Orientale (Italy). His latest book is *Trump e la rivoluzione americana. Da dove viene, dove ci porta* [Trump and the American Revolution: Where It Comes From, Where It Is Taking Us] (Editoriale Scientifica, 2026).

does show, however, that substitution is not built into the technology as its only rational use. (Brynjolfsson, Li and Raymond, 2025).

Broader estimates should be read in the same way. The International Labour Organization calculates that one worker in four is in an occupation exposed to generative AI to some degree, but only 3.3 per cent of global employment falls within the highest exposure category. Because many occupations still require human input, the ILO considers the transformation of tasks more likely than the wholesale disappearance of jobs. Exposure is therefore not a verdict. Between the two lie the organisation of the firm, collective bargaining, public policy and the voice accorded to workers. (Gmyrek et al., 2025).

With the arrival of so-called agents - AI systems capable of receiving a goal and autonomously organising a sequence of operations to achieve it - the stakes rise. A chatbot answers a question; an agent consults records and programs, performs several steps, and finally returns to the person in charge with a result to be checked. A still preliminary study, based on the responses of 1,500 workers across 104 occupations, found that in 45.2 per cent of occupations the preferred level of human involvement was an equal partnership with the system. It also found a gap between what workers would like to automate and the activities on which firms are concentrating investment. This is an important clue: even at the technical frontier of automation, different configurations are possible, and the market does not necessarily select the socially preferable one. (Shao et al., *Future of Work with AI Agents*, 2025).

2. When the Machine Enters the State

As long as the debate remains confined to the firm, politics appears mainly as an external force: it limits the most dangerous uses and supports those who lose their jobs. Artificial intelligence changes this position because the state is also a gigantic productive organisation. Its offices process cases and administer records; the distribution of resources and the exercise of fundamental rights depend on their outcomes. The state is also an employer and a purchaser. Through procurement, incentives and standards, it can shape what firms develop. Once it begins to use systems capable of coordinating complex processes, it no longer merely governs the digital transition from the outside. It becomes one of the places in which the form of that transition is decided.

Public-sector uses are already numerous, even if many remain experimental. The OECD has examined two hundred cases: 57 per cent automate, simplify or personalise processes and services; 45 per cent support decision-making and forecasting; 30 per cent detect anomalies or strengthen accountability. In the British public sector, the Alan Turing Institute estimated that about 41 per cent of working time is devoted to activities that generative AI could support at least in part. Another study by the same institute identified 143 million complex but repetitive public transactions each year and judged 84 per cent of them highly automatable. These are measures of potential, not results already achieved: fragmented data, obsolete IT systems, model errors and organisational resistance may absorb a considerable share of the benefit. But the amount of administrative capacity that could be freed is now too large to be treated as a marginal issue. (OECD, *Governing with Artificial Intelligence*, 2025; Hashem et al., 2025; Straub et al., 2024).

The political question begins here. What will that capacity be used for? The conventional equation of the right with a weak state and the left with a strong state is of little help. AI makes possible different kinds of

BEYOND CREATIVE DESTRUCTION

strong state. A productive-security model uses it to increase national power, compress the cost of civilian functions, strengthen control and discipline the administration. A productive-social model uses it to increase the effectiveness of public action, preserve continuity of employment and devote the productivity it frees to services the state is currently unable to provide. These are ideal types and will overlap in practice. Their purpose is to show that the same word - efficiency - can conceal opposite transformations.

3. AI in the Productive-Security State

The clearest experiment in the first model was the Department of Government Efficiency built around Elon Musk. The executive order of 20 January 2025 tasked it with modernising federal technology and software to "maximize governmental efficiency and productivity". Three weeks later, a second order gave that formula its operational meaning: reduce the workforce and hire no more than one employee for every four employees who departed. Significantly, the ratio did not apply to public safety, immigration enforcement or law enforcement. Technology thus entered the administration together with a political criterion that had already been set. It was not first asked to discover what work was necessary; it was enlisted in shrinking one part of the state. (The White House, Executive Order 14158; The White House, Workforce Optimization Initiative, 11 February 2025).

The reduction was not a metaphor. Data collected by the Government Accountability Office from twenty-two major agencies show that between December 2024 and January 2026 the civilian federal workforce fell by almost 256,000 people, more than 11 per cent. There were 378,000 separations, including incentivised resignations, retirements and forced reductions, while new hires numbered about 127,000. The decline remained around 1 per cent at the Department of Homeland Security and exceeded 45 per cent at the Department of Education. DOGE's temporary mission formally ended on 4 July 2026; a month later, the same GAO found that some of the savings claimed on its "Wall of Receipts" lacked a verifiable method or were simply wrong. The point is not only that the restructuring was conducted chaotically. It is that the number of people removed was treated as a measure of success before any reliable assessment of the public capacity lost. (GAO, Federal Agency Workforce Changes, 17 June 2026; GAO, DOGE Wall of Receipts, 6 August 2026).

DOGE therefore seemed to confirm the popular expectation that AI will arrive in offices to take work away from human beings. Yet the downsizing campaign did not follow from a technical measurement of automatable hours. It was consistent with the long-standing hostility of the American right towards regulatory bureaucracy, social programmes and professional expertise capable of constraining the presidency. AI accelerated the operation and supplied a modern language in which to present an old political project.

Nor did this project amount to a generically smaller state. The 2026 budget request proposed cutting non-defence discretionary spending by 22.6 per cent while increasing defence spending by 13 per cent, to \$1.01 trillion. A further \$175 billion was to go to homeland and border security. As the civilian administration lost staff, security functions were protected from the hiring ratio. The state was not withdrawing: it was shifting capacity and money towards the functions in which the right recognised its mission. (Office of Management and Budget, Fiscal Year 2026 Discretionary Budget Request).

“Right-wing AI” should therefore not be imagined only as the capitalist's robot replacing the worker. Inside the state it can reduce the number of officials who deliver benefits or regulate firms while simultaneously multiplying surveillance capacity and using law-enforcement databases to manage borders. It can also help direct strategic production. The ownership of major technology companies may remain private so long as they accept a role as national infrastructures, protected by government and subordinated to its priorities. The result is a state that becomes lighter in its social and regulatory functions while growing denser around command.

4. AI in the Productive-Social State

Distrust of artificial intelligence arises from real dangers. An algorithm can turn workplace oversight into continuous surveillance and reproduce discrimination buried in data. It can even make a decision effectively unappealable when no one can any longer explain it. The European AI Act and the ongoing consultation on the Quality Jobs Act rightly insist on transparency and human responsibility, protecting workers from excessive monitoring and recognising a role for their representatives. But a policy built only around prohibitions leaves the most important question unanswered: who will use the productive capacity that is being created anyway, and for what purpose? (Regulation (EU) 2024/1689 on artificial intelligence; European Commission, consultation on the Quality Jobs Act, 20 July 2026).

If politics confines itself to opposition, it hands firms and the security state the power to set the goals of innovation. It should begin instead with the distinction between the job and the person. There is no reason to preserve every useless procedure for ever, as though filling in a form by hand were a social right. What must be protected is continuity of employment, income, rights and bargaining power, even when tasks change. The security to be built concerns the worker's trajectory, not the immobility of the office in which that worker happens to sit today.

This reverses the usual sequence. Faced with an organisation that can achieve with two hundred people what once required a thousand, the first question should not be how to tax or redistribute the value produced by eight hundred redundancies. The first task is to determine whether that productivity gain can expand and improve output, or reduce working time while people are trained and redeployed. Only the productivity left after these possibilities have been explored becomes a monetary saving to be divided.

The point is not to give the machine a generic "political prompt". Daron Acemoglu, David Autor and Simon Johnson describe as pro-worker technologies that increase the value of human capabilities, and argue that firms invest too little in this direction. Incentives reward substitution, the dominant ideology equates progress with automation, and the social benefits generated by a better-supported teacher or healthcare worker do not enter in full into the accounts of the company developing the model. Correcting this direction requires public expertise, purposive procurement, tax policy and worker voice. (Acemoglu, Autor and Johnson, Building Pro-Worker Artificial Intelligence, 2026).

The productive-social state should give agents a kind of labour constitution. Higher productivity remains an objective, but it comes after certain constraints. The agent could not propose redundancies while there remained possibilities for redesigning jobs, redeploying staff, financing retraining or reducing working hours. Where rights are at stake, the decision remains human. Any restructuring would also require trade-union participation and an assessment of local and regional effects. In research on multi-objective

optimisation, this is not an exotic idea: lexicographic procedures exist in which the higher-order objective is satisfied first and subsequent objectives are maximised within the space that remains. Operationally, a hierarchy of this kind could be implemented as part of what recent work from the Jeff Bleich Centre describes as a harness: an external, auditable governance layer that imposes policy gates, source rules, response contracts and human-escalation paths across otherwise different models (Kereopa-Yorke et al., 2026; Kereopa-Yorke and Praino, 2026). In the present argument, the lexicographic hierarchy is only one possible component of such a policy harness, not a complete harness in the engineering sense: it constrains what the model is allowed to optimise rather than relying on the model itself to supply the governing priorities. If constraints conflict, the agent does not decide which right to sacrifice. It produces scenarios and returns the choice to political responsibility. (Skalse et al., Lexicographic Multi-Objective Reinforcement Learning, 2022).

Such a mandate acquires force only if the surrounding economic environment changes as well. Training must take place during working hours and be funded. Subsidies and public contracts should depend on credible employment-transition plans, while the tax system should not systematically make it cheaper to buy labour-replacing capital than to invest in people. The OECD finds that consultation and training are associated with better outcomes for workers. In this framework, participation does not arrive at the end merely to negotiate the number of redundancies. It helps decide which tasks should be given to the machine and what new organisation should be built with the time that has been freed. (Lane, Williams and Broecke, 2023).

5. Putting the Framework to Work: Seven Thousand Workers Too Many?

Consider a large Italian public agency with 25,000 employees spread across hundreds of local offices and handling millions of cases each year. Citizens repeatedly have to provide information that the state already holds. Many officials copy data, check forms, compare registers and prepare standard communications. Meanwhile appeals accumulate, difficult cases wait, telephone assistance goes unanswered, and people without adequate digital skills sometimes abandon a benefit to which they are entitled.

An agentic system - that is, one capable of autonomously coordinating a sequence of operations - reconstructs the path taken by cases, identifies the steps that generate errors and distinguishes formal checks from decisions that require judgement. It analyses complaints, waiting times and the territorial distribution of workloads. It can also detect demand that administrative statistics fail to capture: people who abandon a procedure, applications completed incorrectly because the instructions are incomprehensible, substantive checks postponed because offices are overwhelmed by compliance tasks. It concludes that automating repetitive tasks would free an amount of work equivalent to 7,000 full-time employees.

The simplest accounting would translate that result into 7,000 surplus workers. At a conservative total cost of €50,000 per employee per year, the agency could claim a productivity dividend of €350 million. After subtracting €25 million for models, infrastructure, cybersecurity and human oversight, steady-state savings would reach €325 million a year. In Italian public employment, an immediate reduction of this magnitude would not be realistically practicable: the result would instead be pursued through a hiring slowdown, mobility, voluntary departures and a long reduction in headcount. The figure remains useful

BEYOND CREATIVE DESTRUCTION

because it makes the temptation visible. In 2024, the total cost of Italian public-sector employment was €186.1 billion; assuming €50,000 per employee places the example within a plausible and conservative order of magnitude. (Ragioneria Generale dello Stato, OpenBDAP, 2024 data).

The productive-social state gives the same agent a more difficult task. It forbids it to treat hours saved and surplus people as equivalent. It asks the system to search for a configuration in which useless procedures disappear while the capacity released is compared with public needs that remain unmet. The agent must therefore map employees' experience, the adjacent skills they could acquire, their preferences and geographical constraints. At the same time, it estimates how much staff would be needed to clear backlogs, follow a case from beginning to end, prevent errors before they generate appeals, and assist people who cannot navigate digital government on their own.

To test this logic, the constraints just described were submitted to GPT-5.6 Sol, with reasoning set to its maximum level. Without access to real administrative data, the model constructed a scenario exercise: not a forecast, but a possible and internally coherent configuration. The use of a general-purpose LLM here as an exploratory analytical instrument also builds on earlier work in which such models, within specified limits, closely aligned with expert assessments in specialised social-science surveys without task-specific training (Cappelli, Aliberti and Praino, 2024). The task here is different: not to replace an expert, but to explore organisational configurations consistent with explicit constraints.

The resulting configuration allocates the equivalent of 1,500 employees to direct assistance and end-to-end case management; 1,100 to complex cases, appeals and backlogs; 700 to substantive checks that are currently sacrificed to formal verification; and 600 to data quality, security and oversight of algorithmic systems. In total, that is 3,900 units of labour, equivalent to €195 million in administrative capacity redirected towards better services. It is not new expenditure: it is work that already exists and is no longer consumed by the wrong procedure.

A further 1,400 full-time equivalents, worth €70 million, can be converted into time. After 1,700 retirements and voluntary departures, headcount would gradually fall to 23,300. Spread across those employees, the freed capacity would allow a 36-hour week to fall to about 34 hours without any reduction in pay. The 1,700 departures would produce €85 million in gross savings; after the €25 million required to maintain the infrastructure, €60 million a year would remain in the public budget. Integrating records, redesigning procedures and training staff would also require a hypothetical initial investment of €100-150 million spread over three years. The transition would not be cost-free, and the financial benefit would emerge only at steady state.

The final account displays the choice without hiding its cost. Of the €350 million released, €195 million becomes service quality and capacity, €70 million becomes time returned to workers, €25 million keeps the infrastructure under control, and €60 million remains in the public budget. Compared with restructuring based on maximum headcount reduction, the state forgoes monetising €265 million, but that sum does not disappear. It reappears in offices that answer the phone, in cases resolved before they turn into litigation, and in a shorter working week. For prudence, the estimate values services only at the cost of the staff assigned to them and places no monetary value on rights made effective or on citizens' time saved.

BEYOND CREATIVE DESTRUCTION

The most powerful part of the experiment, however, is not the arithmetic division of the dividend. A commission could propose that today. What an agent makes manageable for the first time is the coordination of millions of cases with thousands of professional profiles, across territories where shortages vary and training takes different amounts of time. It can simulate many configurations, update them as new data arrive, and prepare a realistic trajectory for each worker. Someone who knows a procedure well can learn to verify anomalies flagged by the system. Someone with experience dealing with the public can move into case assistance. Automation can be introduced in one office immediately and in another six months later, once training has been completed. Early evidence on the use of agents suggests precisely that their value lies not only in performing a task more quickly, but in enabling organisations to experiment with new workflows at a much lower cost than previous technological transformations. (Johnston et al., *The Shift to Agentic AI: Evidence from Codex*, 2026).

The agent does not invent the new public service on its own, decide which needs deserve priority, or guarantee that every redeployment will succeed. Italian public administrations would first need to make interoperable records that often do not communicate with one another, build internal expertise, and prevent the private supplier from becoming the effective owner of the administrative machine. Decisions affecting rights must remain reasoned, contestable and entrusted to a human decision-maker. In 2026, almost every OECD country already uses AI in at least one area of government, yet only 39 per cent require risk assessments before deployment and just 17 per cent maintain an open algorithm register. Technical capacity is advancing faster than the institutions required to govern it. (OECD, *Digital Government Outlook 2026*).

6. Another Use of Power

AI critics have good reason to fear artificial intelligence owned by a small number of firms and introduced into workplaces with the aim of reducing labour costs. They err when they turn that distrust into a politics of renunciation. Regulating the machine after others have decided its direction means accepting that innovation will be designed elsewhere and arriving only in time to mitigate the damage.

A productive-social state does not promise that no occupation will disappear. It uses its own administration as a laboratory and steers the market through public demand. Collective bargaining brings workers' experience into design. In some cases the state will buy private models; elsewhere it will develop its own infrastructure. In both cases it must require interoperability and the ability to exit contracts, without surrendering the definition of objectives to the supplier. Ownership of the technology matters, but so does who defines the function it must perform and has the power to make that function binding.

Creative destruction describes a dynamic of capitalism, not a constitutional obligation. Efficiency achieved by expelling people is only one way of cashing in the AI dividend. A government can turn that dividend primarily into services and time for life, without giving up a share of financial savings, by acting before job loss becomes the fait accompli around which compensation is then distributed. Artificial intelligence does not thereby acquire a politics of its own. Politics imposes on it a task that the market, left to itself, has no reason to choose: reducing the amount of labour required without producing superfluous human beings.

References

- Acemoglu, Daron, David Autor and Simon Johnson (2026), Building Pro-Worker Artificial Intelligence, The Hamilton Project, Brookings Institution, February.
- Acemoglu, Daron and Pascual Restrepo (2019), "Automation and New Tasks: How Technology Displaces and Reinstates Labor", *Journal of Economic Perspectives*, 33(2), pp. 3-30, DOI: 10.1257/jep.33.2.3.
- Brynjolfsson, Erik (2022), "The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence", *Dædalus*, 151(2), pp. 272-287, DOI: 10.1162/daed_a_01915.
- Brynjolfsson, Erik, Danielle Li and Lindsey Raymond (2025), "Generative AI at Work", *The Quarterly Journal of Economics*, 140(2), pp. 889-942, DOI: 10.1093/qje/qjae044.
- Cappelli, Ottorino, Marco Aliberti and Rodrigo Praino (2024), "The 'Implicit Intelligence' of artificial intelligence. Investigating the potential of large language models in social science research", *Political Research Exchange*, 6(1), Article 2351794, DOI: 10.1080/2474736X.2024.2351794.
- The White House (2025), Establishing and Implementing the President's "Department of Government Efficiency", Executive Order 14158, 20 January.
- The White House (2025), Implementing the President's "Department of Government Efficiency" Workforce Optimization Initiative, Executive Order 14210, 11 February.
- European Commission (2026), "Commission Opens Second-Phase Consultation on Quality Jobs Act", Directorate-General for Employment, Social Affairs and Inclusion, 20 July.
- Government Accountability Office - GAO (2026), DOGE Wall of Receipts: More Transparency Needed on How Savings Are Derived from Contract, Grant, and Lease Terminations, GAO-26-108615, 6 August.
- Government Accountability Office - GAO (2026), Federal Agency Workforce Changes: Update for July 2025 to January 2026, GAO-26-108583, 17 June.
- Gmyrek, Pawel, Janine Berg, Karol Kamiński, Filip Konopczyński, Agnieszka Ładna, Balint Nafradi, Konrad Rosłaniec and Marek Troszyński (2025), Generative AI and Jobs: A Refined Global Index of Occupational Exposure, ILO Working Paper 140, International Labour Organization, Geneva, May, DOI: 10.54394/HETP0387.
- Hashem, Youmna, Jonathan Bright, Sukankana Chakraborty, Kate Onslow, John Francis, Anton Poletaev and Saba Esnaashari (2025), Mapping the Potential: Generative AI and Public Sector Work. Using Time Use Data to Identify Opportunities for AI Adoption in Great Britain's Public Sector, The Alan Turing Institute, June.
- Johnston, Drew, David Holtz, Alex Martin Richmond, Christopher Ong, Prasanna Tambe and Aaron Chatterji (2026), The Shift to Agentic AI: Evidence from Codex, arXiv:2606.26959, 25 June.
- Kereopa-Yorke, Ben, Brigitte Lewis, Laura Schaufelberger, Galina Romantsova, Brian Hill and Corch (2026), The Harness Is the Contract: Architectural Sovereignty for Regulated AI after Robodebt, Jeff

BEYOND CREATIVE DESTRUCTION

Bleich Centre for Democracy and Disruptive Technologies, Flinders University, South Australia, June, DOI: 10.25957/7e64-dh11.

Kereopa-Yorke, Ben and Rodrigo Praino (2026), Sovereign AI Through Governance Harnesses: Evidence from a Multi-Jurisdiction Evaluation, Jeff Bleich Centre for Democracy and Disruptive Technologies, Flinders University, South Australia, August. DOI: 10.25957/79y6-wb39

Lane, Marguerita, Morgan Williams and Stijn Broecke (2023), "The Impact of AI on the Workplace: Main Findings from the OECD AI Surveys of Employers and Workers", OECD Social, Employment and Migration Working Papers, no. 288, OECD Publishing, Paris, DOI: 10.1787/ea0a0fe1-en.

OECD (2025), Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions, OECD Publishing, Paris, DOI: 10.1787/795de142-en.

OECD (2026), Digital Government Outlook 2026: From Foundations to Transformational Impact, OECD Publishing, Paris, DOI: 10.1787/0496b2bc-en.

Office of Management and Budget - OMB (2025), Fiscal Year 2026 Discretionary Budget Request, Executive Office of the President, 2 May.

European Parliament and Council of the European Union (2024), Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence, Official Journal of the European Union, L 2024/1689, 12 July.

Ragioneria Generale dello Stato (2026), Pubblico impiego: dati del Conto annuale 2024 [Public Employment: 2024 Annual Account Data], OpenBDAP, Ministry of Economy and Finance, data updated 1 July 2026.

Schumpeter, Joseph A. (1994 [1942]), Capitalism, Socialism and Democracy, Routledge, London, ch. VII, "The Process of Creative Destruction".

Shao, Yijia, Humishka Zope, Yucheng Jiang, Jiaxin Pei, David Nguyen, Erik Brynjolfsson and Diyi Yang (2025), Future of Work with AI Agents: Auditing Automation and Augmentation Potential across the U.S. Workforce, arXiv:2506.06576, June.

Skalse, Joar, Lewis Hammond, Charlie Griffin and Alessandro Abate (2022), "Lexicographic Multi-Objective Reinforcement Learning", Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence - IJCAI 2022, pp. 3430-3436, DOI: 10.24963/ijcai.2022/476.

Straub, Vincent J., Youmna Hashem, Jonathan Bright, Satyam Bhagwanani, Deborah Morgan, John Francis, Saba Esnaashari and Helen Margetts (2024), AI for Bureaucratic Productivity: Measuring the Potential of AI to Help Automate 143 Million UK Government Transactions, arXiv:2403.14712, March.



**Flinders
University**



Jeff Bleich Centre
for Democracy and
Disruptive Technologies

Contact us
jbc@flinders.edu.au