

The Atrophy of Action: Cognitive Delegation, Digital Interventionism, and the Erosion of Spontaneous Order

Juan Claver

Abstract

The familiar Austrian objection to digital interventionism is aimed at the planner: no central authority can gather the dispersed knowledge it would need to coordinate social life through code. I think the more serious problem lies somewhere else. When the state builds and mandates the infrastructure of modern economic life—programmable money, algorithmic moderation, digital identity—it does not only distort the market process. It wears away the cognitive capacity that the process runs on. Working from Mises’s action axiom, from Kirzner’s account of alertness, and from Hayek’s distinction between outer and inner freedom, I argue that systematic delegation of cognition to state-privileged AI produces what I will call an *atrophy of action*: a slow loss of the capacity for purposeful choice that spontaneous order takes for granted. The upshot is uncomfortable. This kind of intervention is worse than classical socialism, not better disguised—worse, because it throws off no corrective signal and quietly disables the faculties one would need to notice that anything had gone wrong.

A Threat Below the Knowledge Problem

The Austrian tradition makes progress by finding costs that are not yet visible. Mises showed that abolishing markets in capital goods destroys the monetary calculation without which production cannot be rationally ordered [8]. Hayek showed that the knowledge an economy needs is scattered and largely tacit, so that central direction fails for reasons no amount of computing power will ever fix [3]. In both cases the damage is done before anyone can see it, and the economist’s job is to make it visible while there is still time to act.

The standard Austrian line on digital interventionism simply continues this argument. Central bank digital currencies, AI governance regimes, digital identity systems—each concentrates knowledge and decision rights in the state, and so picks up the old calculation and knowledge problems in new clothes. None of this is wrong. But it is incomplete, because it puts the entire defect in the planner. The graver one, I want to argue, shows up in the citizen. When the state builds and compels infrastructure that decides on a person’s behalf—curating which information reaches her, pre-clearing which transactions go through, scoring which counterparties she is allowed to deal with—its main effect is not coercion at all. It makes her own deliberation pointless. And a capacity that is made pointless, systematically and at scale, is a capacity that gets lost.

The capacity in question is the one the whole Austrian account of order sits on: the capacity for purposeful action. So the thesis here is that digital interventionism threatens liberty less by coercing than by atrophy, and that the resulting loss of freedom is one no price registers and no crisis corrects. I defend it on two levels, the a priori and the empirical, and I try to keep the line between them visible the whole way through—because the argument only works if we are careful not to mistake what is deduced for what is merely observed.

Why Cognition Is Constitutive, Not Optional

Praxeology starts from a single proposition: man acts—he applies scarce means to chosen ends [8, 11]. Mises takes this to be apodictically certain, since to deny it is itself a purposeful act. Three things follow from it by deduction alone. They need no empirical support and admit no empirical refutation.

First, valuation is irreducibly first-personal. To act is to prefer one state of affairs over another, and a preference just is the ranking of ends by the subject who acts. An end that some external system ranks *for* an agent is not her preference; it is the system’s. This is not a contingent limitation of today’s technology that tomorrow’s might overcome. It falls straight out of the concept of action, so that any system supplying the valuation displaces the agent from the act as a matter of logic, not of degree.

Second, action runs up against genuine uncertainty. Mises distinguishes *class probability*, where we know how a whole class behaves but not the single case, from *case probability*, where the event is unique and no frequency applies [8, 107]. Real choice is always of the second kind, and from this a clean result about artificial intelligence drops out, provable without reference to any particular machine. A learning system optimises over a distribution drawn from observed cases; it therefore lives, by construction, in the

domain of class probability. It can manage risk. It cannot, in Mises's sense, *act*, because action is defined by its confrontation with the singular and the unprecedented—and no gain in raw capability touches a point that is conceptual rather than technical.

Third, the data of the market are themselves the output of valuation. A price is an exchange ratio that emerges when agents bid according to their own subjective rankings [8, 327–331]. It is not raw information sitting out in the world; it is the residue of human valuing. So if the valuations feeding the price system are really an algorithm's, relayed through nominal acts of choice rather than the agents' own, then what the system aggregates is the designer's judgement and not dispersed individual knowledge. The Hayekian case for the market assumes that individuals are actually valuing. It does not survive a quiet substitution of the valuer.

None of these results leans on an empirical premise. Together they establish that autonomous valuation is not some input the market could get along without: take it away and what is left is not a degraded market but an imitation of one. What they do not show, on their own, is that cognitive delegation *displaces* valuation rather than merely helping it along. That last step is an empirical claim, and Austrian method—which rightly insists that praxeology is a priori [8, 30–41]—forbids me from smuggling it in under cover of the deduction. So let me put it baldly: when an agent habitually hands a cognitive task to an external system, her capacity to do that task unaided declines.

The evidence for this is by now hard to wave away. The well-replicated “Google effect” shows that when people expect to be able to look something up later, they retain less of it themselves [7]. Controlled work finds that the mere presence of a smartphone eats into available cognitive capacity, and that the effect is largest in exactly the people who depend on the device most [10]. Most directly of all, the government-commissioned *International AI Safety Report 2026* records that after three months of AI-assisted work, clinicians' ability to do the same task without the AI had fallen by six per cent, and that across 666 participants heavier use of AI tools went together with weaker critical thinking, the link running through cognitive offloading [5]. There is nothing mysterious about the mechanism. A capacity is kept alive by exercise, and exercise is the one thing delegation removes.

This lets me draw a careful conclusion, and only a careful one. The action axiom does not describe an indestructible essence; it describes a practised capacity. That man acts is certain. How well a particular man, long habituated to delegation, will actually deliberate is a question the empirical record gets to settle. Mises is not contradicted. He is supplemented at precisely the point where praxeology has nothing to say.

Hayek supplies the bridge to liberty. In *The Constitution of Liberty* he separates

freedom-as-absence-of-coercion from *inner freedom*, the capacity to act on one’s own deliberate intentions, and insists that the two together fix “how much use a person can make of his knowledge of opportunities” [4, 15]. The implication gets too little attention in Austrian writing on digital policy, and it is decisive: outer freedom without inner freedom is freedom in name only. An agent who is legally free but deliberatively incapacitated holds her liberty as an empty title. The market needs autonomous valuation; delegation erodes the capacity for it; Hayek names the casualty.

Alertness Cannot Be Delegated

Kirzner identifies entrepreneurial *alertness*—the noticing of opportunities nobody has seen yet—as the force that pushes the market toward coordination [6]. Here, too, the a priori and the empirical have to be held apart. By definition, alertness cannot be delegated to a system trained on the past. What the entrepreneur grasps is an opportunity with no precedent in the data, and that absence of precedent is exactly what separates discovery from calculation [6, 35–38]. A system extrapolating from observed cases can surface patterns already latent in them; the genuinely novel opportunity, by definition, is not in them. So the idea that an AI might be “alert for us” is not so much false as incoherent. It confuses discovery with interpolation.

As a matter of fact—and now we are on empirical ground—alertness is a developed disposition, not a fixed endowment. It is sharpened by practice: by scanning, conjecturing, acting, taking losses, revising. The evidence already cited tells us that pulling out the practice degrades the disposition. An agent who has delegated the scanning for years has not just paused; she has gone without the accumulated, error-corrected judgement that practised alertness actually consists in.

And this gives the result I take to be the most important one here, and the one the tradition least anticipated. The classical knowledge-problem argument says socialism cannot work in principle, but it leaves a way back: restore property and prices, and the entrepreneurial process will, given time, regenerate. Cognitive atrophy shuts that way back. Alertness lost over decades of delegation does not come flooding back the day the mandate is repealed, because rebuilding it takes exactly the autonomous practice the intervening regime had made unnecessary. Put it as a chain. The informational quality of prices depends on the deliberative quality of the valuations going into them; that quality falls under sustained delegation; so a market reopened after a long delegating regime will generate impoverished signals even with its institutions fully restored—not because the mechanism is broken but because its inputs are degraded. The hardware of the market

survives. The human faculty it runs on has wasted away. None of Mises, Hayek, or Kirzner modelled this kind of interventionist damage, for the simple reason that none of them faced a threat to alertness coming not from prohibition but from comfort.

The State as Architect of Cognitive Redundancy

One natural objection is that all this indicts technology, not the state: a freely chosen AI assistant offloads cognition just as effectively as a mandated one. I think state involvement changes the analysis, and on three specifically Austrian grounds.

It forecloses exit. A voluntary market in cognitive assistance is disciplined by competition, so an assistant that leaves its users worse off invites rivals that do not. But where there is one lawful identity layer, one settlement rail, one approved moderation regime, opting out stops being a consumer choice and becomes civic exclusion. And here Mises's theorem that intervention breeds intervention [9] returns in sharpened form, because each new digital layer raises the cost of staying outside it and so manufactures the case for the next.

It suppresses profit-and-loss feedback. In a voluntary setting, the cognitive dependence that harms an agent shows up as worse outcomes she can see and learn from. State infrastructure is engineered, more or less on purpose, to remove that friction: a programmable currency that blocks a category of spending spares the holder the loss that would have taught her financial judgement; a feed that surfaces only sanctioned information spares her the error that would have sharpened her evaluation.

And it destroys the comparative baseline. The voluntary user of an assistant still has her own unaided performance as a yardstick to measure dependence against. When every transaction clears through one rail and every piece of information passes one filter, that yardstick is gone. The loss is then not felt as loss, because nothing uncontaminated is left beside which to feel it. This is a second-order atrophy: the erosion of the very measure by which the first-order atrophy might have been caught.

Three live measures show the pattern in operation. The EU AI Act of 2024 hands a central authority the power to decide which AI applications are permissible, and—perversely—a tool that protects autonomy by laying out conflicting evidence and refusing to resolve the uncertainty for you is easier to cast as “high-risk” than one optimised for frictionless compliance. CBDCs as specified by the major central banks carry the technical capacity for programmable restriction by category, time, or recipient [1]; in praxeological terms this intervenes not in prices but in the act of valuation at the moment of exchange, since money is the medium through which subjective valuations get expressed and com-

pared. And digital identity wallets gate economic participation behind a state-issued credential, retiring the informal, practised judgement of trust in favour of administrative certification—which is to say, retiring the very faculty that judgement of trust exercises.

Worse than Socialism: No Corrective Signal

Austrian theory is clear about why socialism fails, and, just as important, about why its failure is legible: suppress prices and you get misallocation that surfaces as shortage and collapse, and the lesson—restore prices—is there to be learned, however dear the tuition [8, 3]. Cognitive atrophy fails without legibility. There is no price of cognitive capacity, so there is no shortage to mark its decline; no queue outside an empty shop, because the shelves are full and the recommendations are apt. The most the atrophied agent meets is a faint unease when asked to choose for herself unaided—an unease the sheer ubiquity of assistance ensures she rarely has to meet. From the inside the condition is comfortable. The comfort is the danger.

The political economy of it is self-sealing. The faculties one would need to evaluate the interventions that cause the atrophy are the faculties the atrophy degrades. Democratic correction assumes an electorate able to deliberate about its own institutions, and cognitive atrophy weakens that ability while leaving the outward forms of democracy untouched. The dynamic reinforces itself and contains no internal corrective. Socialist economies collapsed, and in collapsing they taught their lesson. A society that has handed its cognitive life over to state-privileged AI need not collapse at all. It may simply persist—comfortable, inert, without the resources required to diagnose its own state.

Huerta de Soto’s concept of *dynamic efficiency* frames the loss precisely [2]. Dynamic efficiency is not the static, Paretian allocation of given means; it is the system’s capacity to throw up new ends, new means, new coordination over time—and that capacity lives in the alertness of acting agents. An economy of cognitively atrophied agents may post perfectly adequate static numbers, its AI intermediaries even nudging measured welfare upward in the near term, while the dynamic capacity on which all later progress depends is quietly hollowed out. By the argument of the preceding sections this is not speculation but consequence: degrade the practised valuation of the agents and you degrade, necessarily, the knowledge-generating power of the order they make up.

Conclusion

Mises located the hidden cost of intervention in the destruction of calculation, Hayek in the suppression of dispersed knowledge, Kirzner in the foreclosure of discovery. Digital interventionism imposes a cost underneath all three: the erosion of the cognitive capacity that makes calculation, knowledge, and discovery possible in the first place. It does not arrive looking like socialism. It arrives as convenience, as safety, as efficiency. It does not forbid autonomous action; it makes autonomous action unnecessary—and what is unnecessary at scale is lost. The standing lesson of the whole tradition is that a free society is not a state of affairs but a practice, kept alive by individuals who deliberate, choose, get things wrong, and discover. Those are cultivated faculties, not guaranteed endowments. A state that, through the infrastructure it builds and compels, makes their exercise needless does not merely intervene in the market. It erodes the human ground the market stands on. To defend liberty under digital interventionism, then, is not only to resist coercion. It is to resist comfort—to insist that the faculty of autonomous choice is not a burden for benevolent systems to take off our hands, but the one irreplaceable condition of a society capable of governing itself.

References

- [1] Bank for International Settlements. Annual economic report 2023. Technical report, Bank for International Settlements, Basel, 2023.
- [2] Jesús Huerta de Soto. The theory of dynamic efficiency. *Procesos de Mercado: Revista Europea de Economía Política*, 1(1):11–72, 2004.
- [3] Friedrich A. Hayek. The use of knowledge in society. *American Economic Review*, 35(4):519–530, 1945.
- [4] Friedrich A. Hayek. *The Constitution of Liberty*. University of Chicago Press, Chicago, 1960.
- [5] International AI Safety Report. International ai safety report 2026. Technical report, Department for Science, Innovation and Technology, 2026.
- [6] Israel M. Kirzner. *Competition and Entrepreneurship*. University of Chicago Press, Chicago, 1973.

- [7] Evan F. Risko and Sam J. Gilbert. Cognitive offloading. *Trends in Cognitive Sciences*, 20(9):676–688, 2016.
- [8] Ludwig von Mises. *Human Action: A Treatise on Economics*. Yale University Press, New Haven, 1949.
- [9] Ludwig von Mises. *Interventionism: An Economic Analysis*. Foundation for Economic Education, Irvington-on-Hudson, NY, 1998. Originally published 1940.
- [10] Adrian F. Ward, Kristen Duke, Ayelet Gneezy, and Maarten W. Bos. Brain drain: The mere presence of one’s own smartphone reduces available cognitive capacity. *Journal of the Association for Consumer Research*, 2(2):140–154, 2017.