Computational Dispossession: Resource Utilization, Infrastructural Dependency, and the Erosion of Independent Productive Capacity

Flyxion
Independent Researcher

November 2026

Abstract

A recent video essay by the channel Non-Compete argues that rising prices for memory and graphics hardware, together with the rapid construction of centralized data centres, should be read as a class phenomenon: an attempt by owners of large-scale infrastructure to withdraw computing capacity from freelancers, gig workers and small businesses and to replace independent ownership with metered access. This essay takes that argument seriously as a structural claim while separating it from the stronger assertion that scarcity is deliberately manufactured. Three results are developed. First, the claim that ownership of a productive instrument confers independence is shown to fail whenever the decisive dependencies of a worker are controlled elsewhere, a point the video makes through the example of ride-hailing drivers and which extends directly to programmers, designers and editors. Second, an engineering argument for pooling scarce memory in centralized facilities, which appears in a clip embedded in the video, is shown to be correct about aggregate utilization and silent about autonomy, and a simple formal framework distinguishes utilization, authorization independence, operational availability and dependency as separate quantities that no general theorem relates monotonically. Third, the historical shift from mainframe terminals to personal computers is used to argue that idle, privately held capacity carries option value that utilization metrics cannot register, and that the more promising responses are engineering efficiency, cooperative governance of pooled equipment, and preserved portability, rather than individual abstention. The essay also identifies a general pattern, the privatization of risk without the decentralization of authority, which explains how nominal independence coexists with infrastructural dependency. A further section, drawing on Catherine Liu's account of the professional managerial class, considers how credentialed administration of access lends moral legitimacy to permission regimes. The argument does not require a conspiracy, and it does not require acceptance of the video's historiography.

Keywords: computational autonomy; infrastructural dependency; resource utilization; means of production; platform labour; option value; cooperative computing; professional managerial class.

1 Introduction

Discussion of the present rise in hardware prices has mostly been conducted in the vocabulary of consumers. Memory modules, storage cards and graphics processors have become expensive, gamers and hobbyists are inconvenienced, and the expected remedy is either patience or a correction in the market for artificial intelligence. A video essay published by the channel Non-Compete, titled *The Hidden Class System Behind the AI Crisis*, proposes a different vocabulary. Its presenter, who identifies himself as EJ, argues that the same price movement is better understood as a labour issue and a question about the means of production, because the people most exposed to it are those who earn their living with their own equipment: freelancers, gig workers, temporary contractors and small business owners in digital fields [13]. The opening minutes of the transcript are explicit about this reframing, and the thesis is stated at about 00:04:02, where the rollout of data centres and the inflation of hardware prices are described as aspects of class warfare through which a capitalist class centralizes means of production in order to compel dependence on facilities it controls.

The argument deserves attention for a reason that is somewhat independent of its political framing. It identifies a distinction that economic discussion of computing frequently omits, namely the difference between owning a productive instrument and purchasing conditional access to one. A person who owns a sufficiently capable computer can begin a computation without asking anyone. A person who rents computation begins each task under terms set by the provider, and those terms can change, lapse or be withdrawn. Whether this difference matters for the livelihoods of independent workers, and whether current developments are enlarging it, are empirical and conceptual questions that can be examined without adopting the presenter's account of capitalist intent.

This essay proceeds in that spirit. Section 2 reconstructs the structure of the video's argument with reference to timestamps in the transcript. Section 3 examines the distinction between ownership and effective control, which is the strongest component of the presentation. Section 4 turns to the utilization argument contained in a clip embedded in the video, and Section 5 identifies the places where the argument moves from structural claims to claims about intention that the transcript does not support. Section 6 gives a formal treatment that separates four quantities which are often conflated. Section 7 considers the historical precedent of the personal computer, Section 8 considers the credentialed administration of access with reference to Catherine Liu's account of the professional managerial class, and Section 9 discusses practical responses. Section 10 concludes.

2 The Structure of the Source Argument

The video is long and discursive, and a considerable portion of it is devoted to explaining the class position of freelancers, which is the topic the presenter regards as a prerequisite for the hardware discussion. Between roughly 00:04:30 and 00:30:00 the presenter distinguishes proletarians, who do not own the means of production, from capitalists, who live on profit alone, and from the petty bourgeoisie,

a transitional class of small owners who hire labour but must also work themselves. The central claim of this middle section is that freelancers do not fit any one of these categories permanently. A temporary worker classified as a contractor may in substance be an ordinary employee without benefits (00:14:56 to 00:15:52), whereas a successful designer with many assistants may be a small capitalist (00:14:35 to 00:14:49), and many individuals oscillate between positions from month to month (00:17:50 to 00:18:21). The presenter treats this instability as a feature of the category and warns against static, moralized classification (00:13:53 to 00:14:33). The timestamps given throughout are those of the auto-generated caption file and are approximate to within a few seconds.

Within this discussion the presenter adopts a criterion that governs the rest of the argument. The question to ask about any worker, he says, is which means of production are essential, who owns those, and whether surplus labour value is being extracted through them (00:27:49 to 00:28:17). Ownership of a peripheral item does not make a worker an owner of the means of production in any meaningful sense, and the examples given are non-slip shoes for a waiter and a laptop or mechanical keyboard for a writer (00:24:27 to 00:25:01).

The hardware argument itself occupies the final third. At about 00:30:11 the video plays a clip, whose speaker is not identified in the transcript, arguing that dispersing the world's limited supply of dynamic memory across personal computers that are used intensively only a small fraction of the time is inefficient compared with pooling it in cloud facilities. The presenter then states at 00:30:58 to 00:31:22 that at least part of what is occurring is an attempt by a ruling class to take away non-essential but still important means of production from freelancers and gig workers. He cites Marx and Engels on the decline of small enterprises (00:31:23 to 00:32:00; 00:35:23 to 00:35:50), calls the process a form of rent seeking and of primitive accumulation (00:32:57 to 00:33:53), and adds a related observation about forward deployed engineers, teams of employees of technology companies who perform integration and troubleshooting work that independent contractors have traditionally done (00:34:16 to 00:35:11). The video closes with a choice between treating a stick of memory as a family relic and joining an organized political response (00:36:39 to 00:37:00).

The structure can therefore be summarized as a chain of four propositions. Data centre expansion increases demand for hardware and thereby raises its price. Higher prices disadvantage independent workers, who own and replace their own equipment. Providers of centralized computation benefit from the dependence that results. And those providers act in order to produce this outcome. The next sections argue that the first three propositions are assessable by ordinary economic evidence and are plausible, while the fourth is not established by anything in the transcript and is not required by the structural point.

3 Ownership and Effective Control

The most interesting passage in the video is its treatment of ride-hailing drivers, which runs from approximately 00:19:13 to 00:24:03. A common view holds that a driver who owns a car owns a means of

production and is therefore petty bourgeois. The presenter replies that this confuses phenomena with essence. The car is a means of production, but it is not the one that generates the decisive income. The platform, which determines how customers are matched, how prices are set and how payment is distributed, is the essential means of production, and the driver has no ownership of or control over it (00:22:27 to 00:22:56). The driver, in the presenter's account, bears the depreciation, maintenance and operating costs of the physical asset while the company retains control over the infrastructure that makes the asset economically productive (00:23:12 to 00:23:46). Empirical work on driver experience is consistent with this description in its emphasis on information asymmetry and algorithmic management, both of which place the control of pricing and assignment on the platform side [15].

The point generalizes into a criterion that is more useful than the class labels it was introduced to settle. Ownership of an individual productive asset does not establish economic independence. What establishes it is control over the whole set of dependencies required to turn labour into income. A worker may own a vehicle, a camera, a workstation or a laptop and still depend on a platform for customers, on a vendor for software, on a payment intermediary for settlement and on a distribution channel for audience. The presenter lists several such cases for creative workers, including subscription pricing for graphics software and the commissions taken by freelancing marketplaces and video platforms (00:25:37 to 00:26:16).

The same reasoning applies to computational work. A programmer who owns a capable machine may nonetheless depend on a proprietary development environment, a cloud build service, a model interface offered through an application programming interface, an app store or marketplace through which software is distributed, and a payment processor. Each of these is a point at which the worker can be excluded, repriced or altered. The economic literature on platforms has described this arrangement from several directions, for example as a shift of control from producers of goods to the owners of the intermediating infrastructure [16], or as a form of rent collected by those who control access [18]. What the video adds is a specific proposal, that the hardware itself is becoming another such point, which is addressed below.

Long-term contract work makes the structure especially clear, and it suggests that three forms of productive capacity should be kept apart: possessing the skills needed to perform work, possessing the physical instruments needed to perform it, and possessing the commercial relationships through which it is compensated. A worker can control the first two almost completely and remain dependent on another party for the third. A dispatch or staffing intermediary that controls the relationships with customers controls the conversion of skill into income, however broad the skill and however complete the equipment. There is also an asymmetry in contracting over many years. The intermediary can benefit from accumulated experience, reliability and institutional knowledge without assuming the long-term obligations of an employer, so the worker experiences continuity of labour without continuity of economic security. Whether such an arrangement is legally employment is a separate question from how it is designated by contract, and the video's discussion of temporary workers treated as contractors on paper makes the same point (00:14:56 to 00:15:52).

A general pattern follows, which may be called the privatization of risk without the decentralization of authority. Workers are treated as independent enterprises when costs and liabilities must be allocated, so that depreciation, maintenance, tools, insurance and the consequences of damage fall on them, and as dependent participants when access, compensation and operating conditions must be determined, so that these are set by the intermediary or platform. The ride-hailing example has exactly this form (00:23:12 to 00:23:46). The pattern explains how nominal independence coexists with profound infrastructural dependency, and it supplies a diagnostic that does not depend on class labels: compare the share of risk and cost a worker bears with the share of decisive authority she holds. A large gap between the two marks a position in which independence is a legal description and not an economic fact. In the notation of Section 6, it is a case of high cost burden together with low authorization independence A_i and high dependency D_i.

It is worth noticing an internal tension in the presenter's own framework here. His criterion makes the laptop an unimportant means of production, as the passage at 00:24:27 to 00:25:01 insists, and he designates the computers whose withdrawal he worries about as non-essential (00:04:15; 00:31:08 to 00:31:17; 00:32:57 to 00:33:10). Yet the argument that hardware scarcity constitutes primitive accumulation requires that this equipment matter a great deal for the independence of the workers concerned. The resolution is available but is not given in the video. A peripheral instrument may be non-essential in the sense that its ownership does not by itself generate profit, and yet be important in the sense that its absence forces a worker into the dependence that the essential platforms already impose. Owning the computer is then valuable not because it is a source of surplus but because it is an exit option, a point developed in Section 7.

4 The Utilization Argument

The clip embedded at 00:30:11 to 00:30:44 states an engineering position in a compact form. Many people hope that open models will remove the need for data centres because the models can be run locally, but from the standpoint of global allocation it is poor practice for every individual to hold a share of the world's limited dynamic memory that is in use only a small fraction of the time, and a supply-constrained environment calls for maximum utilization in the cloud. The figure of five per cent is given in the clip without a source and should be treated as an illustrative assertion.

The position is correct on its own terms. If the objective is to minimize the quantity of hardware needed to serve a given pattern of demand, pooling is more efficient than dispersal, and the reason is the familiar one of statistical multiplexing. Suppose n users each demand one unit of capacity at any instant with probability p. Provisioning each user locally so that demand is always met requires n units, and the resulting utilization is p. Provisioning a shared pool so that the probability of shortfall does not exceed a tolerance ε requires, when demands are independent and n is large enough for a normal approximation, about $C_n \approx np + z_\varepsilon \sqrt{np(1-p)}$ units, where z_ε is the standard normal quantile at level $1-\varepsilon$. Capacity is in practice provisioned in integer units, so this is an approximation. Taking offered

load as a proxy for served load, the utilization of the pool is then

$$U_{\text{pool}} \approx \frac{np}{C_n} = \frac{1}{1+z_\varepsilon\sqrt{\dfrac{1-p}{np}}}.$$

This tends to one as n grows, and the capacity needed per user falls toward p. The independence assumption is consequential, however. Exchangeability and a pairwise correlation do not by themselves justify a normal approximation when the correlation stays positive as n grows, so a latent-variable formulation is more appropriate. Let $\Theta \in [0,1]$ be a latent demand intensity with mean p, and let users be conditionally independent given Θ, so that $X_i \mid \Theta \sim \text{Bernoulli}(\Theta)$. Then $\text{Var}(\Theta) = \rho\, p(1-p)$, where ρ is the pairwise correlation of the demands, and by the conditional law of large numbers $\frac{1}{n}\sum_{i=1}^{n} X_i \to \Theta$ almost surely. Under mild quantile regularity conditions the asymptotic capacity required per user is

$$\lim_{n\to\infty} \frac{C_n}{n} = F_\Theta^{-1}(1-\varepsilon),$$

the $(1-\varepsilon)$ quantile of the latent intensity, and the limiting offered-load utilization is $p/F_\Theta^{-1}(1-\varepsilon)$. This depends on the whole distribution of Θ and not on its mean and variance alone. A Gaussian approximation to the quantile gives $p + z_\varepsilon\sqrt{\rho\, p(1-p)}$, but synchronized demand can be extreme, discontinuous or multimodal, as when a popular model is released, so that figure is indicative only. The correlation parameter is not a technicality. Computational demand is synchronized by working hours, software releases, market events, academic deadlines and the arrival of new models, so the assumption of uncorrelated users should be examined and not presumed. The expression also uses offered load, whereas realized utilization depends on served load, which is smaller because a pool occasionally rejects or defers work. A designer who treats hardware as the scarce quantity still has good reason to favour the pool, but the size of the advantage depends on the correlation structure of demand.

The difficulty is that the argument assumes the scarce quantity is the only quantity worth measuring. From the standpoint of an operator, idle capacity is an inefficient use of capital. From the standpoint of an independent owner, the same idle capacity is availability, resilience and the ability to begin a project that was not anticipated. The calculation above says nothing about who has the authority to start a computation, under what conditions, and with what guarantee of continuation. It also rests on a premise that is partly produced by the thing it justifies. The scarcity that makes pooling attractive is plausibly in part a result of the demand from the very facilities for which pooling is recommended, though the extent of this is an empirical matter, so the argument has a circular element: concentrated demand creates shortage, and shortage is cited as the reason that capacity must be concentrated. This does not make the argument false, since a shortage is a shortage whatever its cause, but it does mean that the efficiency case and the dependency case are not independent of each other, and that a consumer who feels the price effect is entitled to question the framing.

5 Where the Argument Outruns Its Evidence

The video repeatedly moves from a defensible structural claim to a stronger claim about deliberate strategy. At 00:01:41 to 00:01:54 the ruling class is described as spamming data centres across the globe. At 00:04:09 to 00:04:29 the price inflation and rollout are called an attempt by the capitalist class to centralize means of production coercively. At 00:33:19 to 00:33:53 the process is called primitive accumulation in which technology companies forcibly and unilaterally seize means of production by raising costs.

It is useful to separate four propositions. The first is that data centre construction raises demand for memory and related components. The second is that scarcity of these components disadvantages independent workers relative to firms with procurement power. The third is that providers of cloud computation benefit from increased dependence among workers who cannot afford equipment. The fourth is that providers intentionally manufacture scarcity in order to dispossess independent workers. The first three could be established with evidence on component markets, procurement contracts, price series and the financial position of independent workers. Even the first is an empirical market claim that should not be assumed, since consumer memory, server memory and high-bandwidth memory are related but not interchangeable products, and the degree to which data centre demand drives the price of consumer parts depends on shared fabrication capacity and on allocation decisions by manufacturers. The fourth would require evidence of corporate strategy, coordinated purchasing aimed at foreclosing rivals, or manipulation of supply, and the transcript offers none. The presenter himself concedes that the present moment might be an ordinary market correction and that his thesis could be obsolete (00:03:09 to 00:03:59), and the claim about wages for freelancers is attributed only to an unnamed study from the previous year (00:06:09 to 00:06:21). These hedges are honest, but they leave the strong claim unsupported.

The usage of the term primitive accumulation also deserves caution. In the classical account the term refers to the historical separation of producers from their means of subsistence through enclosure, conquest and other forms of extra-economic force [12]. Higher prices produced by market competition for a scarce component are not obviously of this kind, since nobody is legally barred from buying the component and no producer is expelled from land. A better description might be the ordinary operation of capital accumulation, in which rising capital requirements raise barriers to entry, an effect the presenter correctly associates with the passages from Marx and Engels that he reads [11]. The structural argument needs nothing stronger than this.

Crucially, the structural argument does not require a conspiracy. Individually rational commercial decisions can collectively generate a system of dependency that none of the participating organizations explicitly designed. A firm that purchases memory at scale because its customers demand inference, a second firm that does the same, and a manufacturer that allocates capacity to the buyers who pay the most are each behaving ordinarily. Their combined effect may nonetheless be to raise the price of the same component for independent owners and to make the rented alternative relatively more attractive.

Whether this is described as dispossession depends on what is thought to have been possessed, and Section 7 argues that the answer is a capacity for unmediated initiative.

The presentation also treats the displacement of independent producers as nearly inevitable, in line with the cited statement that the petty bourgeoisie is a dying breed on a long enough timeline (00:29:35 to 00:29:58). This neglects several countervailing developments. Open-source software and publicly documented standards reduce dependence on proprietary tools. Commodity hardware and second-hand markets extend the useful life of equipment. Distributed and cooperative arrangements allow capacity to be pooled without being owned by a single firm, a possibility the commons literature has long treated as viable under suitable governance [14]. And improvements in computational efficiency can lower the minimum capital needed to participate. None of these guarantees an outcome, but together they make inevitability an overstatement.

6 A Formal Separation of Utilization, Autonomy and Dependency

The argument so far suggests that several quantities are frequently run together, and that precision requires distinguishing four of them. The first is the utilization of a body of infrastructure,

$$U \;=\; \frac{\text{actively consumed capacity}}{\text{installed capacity}}.$$

The second is the authorization independence of a worker i. Let T_i be the tasks she requires, each with weight w_t, let V be the set of relevant resources, services and permissions, and let $L_i \subseteq V$ be those under her independent control. Each task t has a family $\mathscr{S}_i(t) \subseteq 2^V$ of sufficient configurations, any one of which allows the task to be executed. Then

$$A_i \;=\; \frac{\sum_{t \in T_i} w_t \, \mathbf{1}\big[\exists S \in \mathscr{S}_i(t) :\; S \subseteq L_i\big]}{\sum_{t \in T_i} w_t}.$$

This formulation captures substitutability and redundancy, and it explains why open standards and interchangeable components matter economically: they enlarge the family of sufficient configurations and so reduce the chance that one provider can obstruct a whole workflow. The third quantity is operational availability, R_i, the weighted fraction of initiation attempts that are actually served within a tolerable delay. It is distinct from authorization, because an enforceable right of access does not remove congestion, equipment failure or scheduling conflict. The fourth is infrastructural dependency, which is better described by a withdrawal susceptibility profile than by a single number. For $k \geq 0$ let

$$D_i(k) \;=\; \max_{\substack{W \subseteq V \setminus L_i \\ |W| \leq k}} \frac{\sum_{t \in T_i} w_t \, \mathbf{1}\big[\forall S \in \mathscr{S}_i(t) :\; S \cap W \neq \varnothing\big]}{\sum_{t \in T_i} w_t},$$

the largest weighted share of her tasks that can be disabled by withdrawing at most k external resources, and write $D_i = D_i(1)$ for the single-resource measure. Assuming every task initially has a sufficient

configuration, $D_i(0) = 0$ and $D_i(k)$ is nondecreasing in k, and when all external resources are withdrawn $D_i(|V \setminus L_i|) = 1 - A_i$, which ties dependency exactly to authorization independence. The profile shows why a single-resource measure can mislead. A worker who can run a task on either of two competing providers has $D_i(1) = 0$, yet if neither is under her control she has $A_i = 0$: she is protected against the failure of one provider while wholly dependent on an externally controlled industry. For institutional analysis the withdrawal sets should eventually be defined over controlling authorities and not over individual resources, since two apparently interchangeable services may belong to one corporation, rely on one identity provider or fall under one jurisdiction, and redundancy of resources is not redundancy of authority. A simple count of dependencies would be misleading, since it changes when a dependency is merely decomposed into parts, and a worker who controls ninety-nine of a hundred inputs can still be incapacitated by the withdrawal of the last. Utilization is a property of a facility or fleet. The other three are properties of a worker's relation to the facilities she uses, and they depend on governance as well as on hardware. The distinction between a task being technically possible and a worker having standing to initiate it is the point of the autonomy measure, since a task executed under a revocable permission is not independently executable in the intended sense.

Proposition 1 (Pooling and utilization)**.** *Let n users have independent demands, each active with probability p. Under the normal provisioning approximation, the offered-load utilization of a pool provisioned to a shortfall tolerance ε tends to one as n grows, and the capacity required per user tends to p. If instead demands are conditionally independent given a latent intensity Θ with $\mathrm{Var}(\Theta) > 0$, the capacity required per user tends to $F_\Theta^{-1}(1-\varepsilon)$ and the limiting offered-load utilization to $p/F_\Theta^{-1}(1-\varepsilon)$, which is below one whenever that quantile exceeds the mean p, as it does for all sufficiently small ε.*

This is the claim contained in the embedded clip, and it is accepted subject to the stated demand model and provisioning approximation. The following proposition states the limitation.

Proposition 2 (Non-monotonicity of autonomy in utilization)**.** *There exist allocations serving the same population of workers under the same offered demand process, with possibly different quantities of installed hardware, such that an increase in utilization is accompanied by a decrease in A_i for every worker, by no change in A_i, and by an increase in A_i. There also exist two physically identical facilities with identical utilization and different values of A_i. Consequently authorization independence is not a monotone function of utilization, and it cannot be identified from physical infrastructure metrics alone.*

Proof. Let x be local ownership in which each of n workers holds one unit and all tasks of each worker are executable on that unit without permission. Then $U(x) = p < 1$ and $A_i(x) = 1$. Let x' be a pool provisioned as in Section 4, operated by a single provider that grants or withholds each execution at its discretion, with no worker holding an enforceable guarantee of access, and that in ordinary operation serves admitted demand, so that its possession of a veto does not lower realized utilization. Then $U(x')$ exceeds p by Proposition 1, but every execution requires authorization, so no task is independently executable and $A_i(x') = 0$ for every i. For the second case, let x'' be the same physical pool owned and

governed cooperatively by its users under a charter that guarantees each member access up to a stated quota and prevents any external party from vetoing an execution. Then $U(x'') = U(x') > U(x)$ and $A_i(x'') = 1$, because no authorization external to the workers is required, although $R_i(x'')$ may fall short of what a commercial provider offers, since quotas do not remove congestion, failure or scheduling conflict. For the third case, let y be an arrangement in which each worker rents an individually allocated, underutilized machine from a provider with discretionary control over access, so that $U(y) = p$ and $A_i(y) = 0$, and replace it with the cooperative pool x''. Then $U(x'') > U(y)$ and $A_i(x'') = 1 > A_i(y)$. The three constructions together show that increasing utilization can accompany decreasing, unchanged or increasing autonomy. Finally, x' and x'' are physically identical and, under the stipulation above, have identical utilization, yet $A_i(x') = 0$ and $A_i(x'') = 1$, so autonomy cannot be recovered from utilization or from capacity alone. □

The proposition is elementary, but it removes a common source of confusion. When it is said that pooling is efficient, the statement concerns U. When it is said that pooling is dangerous, the statement concerns A_i and D_i. The two statements are compatible, and the question of which should govern allocation is a question of values and institutions, not of arithmetic. The cooperative case in the proof also shows that the problem is not physical centralization. The relevant variable is who holds the authority to initiate computation, a variable that centralized machinery does not by itself determine. Authorization independence must also be kept apart from operational availability. A cooperative can offer full independence with imperfect availability, and a commercial provider can offer excellent availability while retaining unilateral power to terminate access, so A_i and R_i are different properties and neither implies the other.

A further consideration concerns tasks that cannot be forecast. Let unanticipated tasks arrive at rate λ with value $v \geq 0$. Let c_L be the marginal cost of local execution, π the price of remote execution, q the probability that remote execution is both available and authorized, and K_L the amortized cost per unit time of maintaining local capacity. Because a task is undertaken only if its value exceeds its cost, the expected value flows are $V_L = \lambda\, \mathbb{E}[(v - c_L)_+]$ for local execution and $V_C = \lambda q\, \mathbb{E}[(v - \pi)_+]$ for remote execution, and the net advantage of local capacity is

$$\Delta V \;=\; \lambda \Big(\mathbb{E}[(v - c_L)_+] - q\, \mathbb{E}[(v - \pi)_+] \Big) - K_L .$$

This expression can take either sign, so the argument does not presuppose that ownership is always preferable. It does show what drives the sign: a low q, a high π and a high rate of unanticipated tasks favour local capacity, whereas a high K_L favours the provider. None of this is visible to a metric that counts only realized use. The more interesting extension is to let the probability of authorization depend on the task and the requester, writing $q(t, i)$ in place of a constant. Selective exclusion, provider restrictions on unconventional research and commercial prioritization then become explicit terms, and the option value of independence is largest for exactly the speculative, amateur and experimental projects whose value is not yet recognized. This is the standard logic of valuing flexibility under uncertainty

[4], applied to computing capacity.

The practically important case is a hybrid, in which a worker uses commercial services while keeping local capacity as a fallback. Let $Z \in \{0,1\}$ indicate whether remote execution is available and authorized, with $\Pr(Z=1)=q$. For tasks that either system can execute equivalently, a worker holding both obtains

$$V_H = \lambda\,\mathbb{E}\big[\max\{0,\ v-c_L,\ Z(v-\pi)\}\big] - K_L,$$

whereas the remote-only arrangement yields $V_C = \lambda\,\mathbb{E}\big[Z(v-\pi)_+\big]$, so the value of the fallback is

$$V_H - V_C = \lambda\,\mathbb{E}\big[\max\{0,\ v-c_L,\ Z(v-\pi)\} - Z(v-\pi)_+\big] - K_L.$$

The local machine may sit idle precisely because the commercial service is currently attractive, and its idleness does not show that it is redundant. The value of insurance is not measured by how often a claim is made, and local capacity insures against interruption, exclusion, price increases and changes in contractual terms.

7 The Personal Computer as Precedent

The historical analogy with the mainframe period clarifies what is at stake. What made personal computing consequential was not simply that machines became smaller and cheaper. It was that the authority to initiate computation passed from institutions to individuals. Before the hobbyist movement, serious computation depended on employment, university affiliation, institutional approval or the purchase of time on another party's machine. The hobbyist movement, with its clubs, technical publications, shared software and documented hardware designs, allowed an individual to run an experiment without persuading an institution that it was worthwhile [9, 17].

Seen this way, the contemporary equivalent of the mainframe terminal is an inexpensive consumer device attached to a very expensive computational facility over which its user has little control. The analogy is imperfect in an instructive way. Institutional computing also mediated access to data, software, research opportunity and professional resources, so the difference is not that the mainframe controlled less. It is that the breadth and pervasiveness of present dependencies is greater, since a single provider relationship can now extend across software, storage, distribution, identity and payment. The history of computing is better read as an oscillation between arrangements of execution authority than as a progression. Institutional mainframes placed authority with administrators, commercial timesharing with service providers, hobbyist microcomputers and the personal computer era with individual owners, networked computing with individuals and network operators jointly, and cloud and subscription software with infrastructure and software providers, while local open-weight models point toward individual or cooperative operators once more. These arrangements coexist and do not wholly replace one another, but the pattern shows that technological progress does not by itself produce a monotonic increase in individual autonomy. Computational recentralization is a suitable name for the present

tendency, and it should be distinguished from technological regression: the machines become more powerful while the individual becomes less able to exercise independent control over them. Aggregate capability and resource efficiency can both rise while independent productive capacity falls, and there is no contradiction among the three.

The same history supplies the philosophical reply to the utilization argument. A public library is not inefficient because its books are not all being read at once, and a workshop is not wasteful because its tools are not all in use. Idle capacity is a precondition of spontaneous activity. If every computational resource is allocated according to immediate and measurable demand, speculative work must compete for resources against established commercial uses, and the allocation mechanism will consequently favour activities whose value is already recognized over those whose value remains to be discovered. Much of the history of computing consists of people finding uses for machines that their makers did not anticipate. The achievement of the personal computer was to give ordinary individuals a small domain of unmetered experimental freedom, and the preservation of that domain may matter as much to innovation as the growth of aggregate capacity.

The presenter's opening remarks, in which a collection of computers is described as the nearest thing to intergenerational wealth he has owned (00:00:56 to 00:01:32), point at the same value from the household side. A computer held in reserve is a store of capability. Its worth lies partly in what it permits without anyone's consent, and this is why a price rise for such equipment is felt as more than an inconvenience.

The personal computer has a further significance for workers in precarious or unrelated occupations. Historically it allowed a person to finish a day of physical or routine labour, return home and undertake software development, scientific research, writing or experimentation on equipment owned outright, without needing another employer, dispatcher or intermediary to authorize the activity. It was therefore not only a productive instrument but a means of escaping dependence on one's existing occupational circumstances. If advanced computational work comes to require subscriptions, metered inference, rented storage and access to proprietary infrastructure, that avenue acquires the structural characteristics of contract labour: the individual bears the costs of developing expertise, maintaining equipment and undertaking speculative work, while another party controls the conditions of participation. This is the same privatization of risk without decentralization of authority identified in Section 3, now applied to a person's second occupation. It is also why the preservation of independently owned computational capacity matters to people whose primary work has little to do with computing.

8 Credentialed Administration of Access

A further strand of the argument concerns who administers a permission regime once one exists. Catherine Liu's *Virtue Hoarders* offers a vocabulary for this question, and the vocabulary is useful provided that its status is kept in view. The book describes itself as a polemic, and its generalizations about the professional managerial class should be read as claims to be tested and not as findings. Liu,

building on the concept introduced by Barbara and John Ehrenreich [5], treats this class as a stratum of credentialed salaried workers in academia, the professions, the media, foundations, government and corporate management. Her thesis is that the class serves capitalist arrangements while presenting itself as progressive, converts taste, consumption and moral display into status, and promotes meritocracy and philanthropy as individualist substitutes for redistribution [10].

The connection to the present topic begins with a passage the video itself reads. In the Manifesto the small independent owners are described as being replaced in manufacture, agriculture and commerce by overlookers, bailiffs and shopmen (00:29:28 to 00:29:34), that is, by a managerial stratum acting for the owners of concentrated capital. The video turns to the forward deployed engineers of technology firms at 00:34:16 to 00:35:11 as a contemporary instance. In a regime where computation is obtained by authorization, someone must administer the authorizing: those who write usage policy, review requests, screen applicants, allocate grants, set quotas and decide which uses count as responsible. The proposal here is that this administrative layer is staffed largely by credentialed professionals whose standing rests on certification and on claims to virtue such as safety, responsibility and ethical stewardship, and that this shapes the permission regime in two ways.

The first effect is on the authorization probability introduced in Section 6. If $q(t, i)$ depends on the requester as well as the task, then conditioning access on credentials or institutional affiliation lowers q for exactly those workers who lack them, namely independent researchers, self-taught programmers and freelancers. Some programmes that grant access to restricted capability or subsidized compute condition eligibility on institutional affiliation and others do not, and how common the practice is, and how much of it is legitimate eligibility screening and how much is credential-based exclusion, is an empirical question this essay does not settle. Where such conditions exist, they are a reasonable screening device for an administrator and a substantial barrier for someone outside the institutions that issue the credential. The independence measure A_i is then lower for the uncredentialed even when hardware, price and technical skill are held fixed.

The second effect is on contestability. When access is restricted in the name of safety or responsible stewardship, an objection to the restriction can be presented as an objection to safety. The utilization argument in the embedded clip has a similar rhetorical structure, since maximum utilization is offered as a form of responsible stewardship of a scarce resource and the opposing intuition is described as mistaken. A moral rationale of this kind converts a distributional question, who may initiate computation, into a question of conscientiousness, and this is a form of what Liu calls virtue hoarding: the accumulation of moral standing alongside the control of the scarce good. The parallel is structural, not literal, since Liu's concern is the class politics of taste and consumption and not compute governance.

The structure involves three positions that should not be conflated. The infrastructure owner holds ownership, investment and residual financial claims. The credentialed administrator holds delegated authority to authorize, classify and exclude, and need not own any of the equipment, being in many cases a salaried employee who is herself dependent on the owner. The independent requester holds only conditional execution and her own labour. The administrator's authority is neither ownership nor

ordinary technical expertise. This matters because a permission regime can expand without any further concentration of capital: an institution may add screening procedures, certification requirements, eligibility criteria or discretionary reviews while keeping exactly the same equipment. There are therefore at least two independent mechanisms of computational dispossession, material exclusion through rising resource requirements and administrative exclusion through increasingly restrictive conditions of access, and the second can operate even as computation becomes cheaper. It corresponds to a form of scarcity distinct from physical and economic scarcity, namely administrative scarcity, in which resources exist and may be affordable but their use requires authorization that can be withheld.

Credential-based allocation also has a circularity that requires no ill will to operate. Suppose access to advanced resources depends on demonstrated accomplishment, while producing accomplishment requires such access. Let $h_i(t)$ be an individual's accumulated recognized standing, let $q_i(t) = g(h_i(t))$ be her probability of access with g increasing, and let standing evolve as

$$\frac{dh_i}{dt} = \alpha_i q_i(t) - \delta h_i(t),$$

where α_i is her potential productivity and δ the depreciation of recognized accomplishment. Equilibria satisfy $\alpha_i g(h) = \delta h$, and an equilibrium is stable when $\alpha_i g'(h) < \delta$. If g is sufficiently nonlinear, for instance S-shaped, then for intermediate values of α_i this equation has three solutions, of which the lowest and highest are stable. Two individuals with identical productive potential can then settle at different equilibria depending only on initial standing, and the low equilibrium is one in which a capable person cannot acquire the accomplishments needed to qualify for the resources needed to demonstrate capability. This is a stylized model and not an empirical finding, but it shows how ostensibly meritocratic rules can reproduce the conditions they purport to measure, without malicious administrators or deliberate discrimination. It gives a formal counterpart to Liu's criticism of meritocracy that does not rest on imputing motives. The relevant empirical question is whether an allocation mechanism measures independent competence or competence already augmented by privileged access.

Contestability should also enter the framework. Two permission regimes can have the same authorization probability while differing in institutional character. One may reject a request according to published rules, give a specific reason, allow an appeal and offer an alternative execution pathway, while another rejects the same request without explanation or recourse. Contestability concerns whether an adverse decision can be examined, challenged and corrected through a procedure that is not wholly at the discretion of the deciding authority, and it can be summarized by the probability c_i that a wrongly denied request is reversed on independent review. The effective probability that a legitimate request succeeds is then $q_i + (1-q_i)c_i$, so regimes with equal q_i can leave the requester in very different positions.

Three cautions apply. Some restrictions on powerful capabilities are well motivated, and some credential screening serves a legitimate function in limiting misuse, so the argument does not show that every such practice is pretextual. Liu's thesis has the same evidential difficulty as the video's account

of capitalist intent, since it explains by class interest what institutional incentives, risk aversion and administrative convenience might explain equally well, and Section 5 applies equally here: the structural point requires no assumption about motive. Finally, the independent worker is not simply opposed to the credentialed stratum, because many freelancers and independent researchers hold credentials themselves, and the relevant line runs between those who administer access and those who must request it. What the framework contributes is a reason to ask, of any proposed permission regime, who staffs it, on what grounds its authority is justified, and whether those grounds can be contested by the people it excludes.

9 Responses

If the diagnosis is correct, the appropriate responses operate at several levels, and none of them depends on the political program with which the video concludes. The first is engineering. Making useful computation require fewer resources lowers the capital needed to take part in computational work, and in this respect techniques such as quantization, distillation and sparse or memory-efficient methods are as much economic as technical developments [7, 6, 3]. A capability that formerly required a large installation and can be reproduced on modest hardware has had its political economy altered without any change in law and without confrontation with any firm. There is a necessary qualification, which is that efficiency gains must remain accessible to independent users. If every improvement is absorbed into larger proprietary systems, efficiency reinforces concentration instead of relieving it.

Efficiency alone is not sufficient, and the point can be stated formally. Suppose the resources needed for a capability of quality Q at time t are $C(Q,t) = f(Q)/\eta(t)$, where η measures computational efficiency, and let $B_i(t)$ be the resources an independent researcher can afford and $Q_*(t)$ the prevailing standard of competitive performance. Independent participation is feasible when $C(Q_*(t),t) \le B_i(t)$. Taking logarithmic derivatives of the ratio gives

$$\frac{d}{dt}\log\frac{C(Q_*(t),t)}{B_i(t)} \;=\; \varepsilon_f(Q_*)\,\frac{\dot{Q}_*}{Q_*} \;-\; \frac{\dot{\eta}}{\eta} \;-\; \frac{\dot{B}_i}{B_i},$$

where $\varepsilon_f(Q) = Qf'(Q)/f(Q)$ is the elasticity of cost with respect to quality, so participation deteriorates whenever the growth rate of the competitive standard, weighted by that elasticity, outpaces the combined growth of efficiency and affordable resources. Independent researchers can therefore own increasingly powerful machines while becoming less competitive, because absolute progress can coexist with relative exclusion. Democratization requires that efficiency and affordable capacity keep pace with the escalating requirements of participation.

The second response is cooperative ownership. A group of independent researchers or small firms who cannot individually afford an expensive workstation may collectively maintain several. The decisive difference between such an arrangement and a commercial cloud is not necessarily the physical location of the equipment but the governance of it, as the proof of Proposition 2 shows. A cooperative

facility can capture much of the multiplexing gain of Proposition 1 while leaving authority over initiation with its members. The record of long-lived commons suggests that this is difficult but feasible when membership, quotas and enforcement are clearly specified [14, 1]. Collective ownership is not by itself autonomy-preserving, however. Administrators of a cooperative can become entrenched, membership requirements can become exclusionary and internal procedures opaque. The relevant protections are enforceable execution rights, transparent allocation rules, meaningful participation, independent review and practical exit, so cooperative ownership is a possible means of securing these properties and not a guarantee of them.

The third response is the preservation of exit. It is possible to use commercial infrastructure heavily without becoming irreversibly dependent on it, provided that source code, data, documentation and executable workflows are kept in local copies, formats are open, environments are reproducible and the interfaces to services can be exchanged. Distributed version control is a useful model, since a repository on a corporate host is still a portable object whose contents and history can be replicated elsewhere, so the dependency remains reversible as long as independent copies are maintained. The notion of exit and its relation to voice and loyalty is the subject of a well-known analysis [8], and its practical implication here is that the ability to leave is itself a form of bargaining power.

The fourth response concerns individual consumer abstention, which is unlikely to be sufficient. Large firms possess procurement advantages, economies of scale and access to capital that personal frugality cannot reproduce. Durable responses therefore combine engineering alternatives, independent infrastructure, cooperative arrangements, interoperability and, where appropriate, institutional measures concerning competition, repairability, portability and public research computing. The objective need not be to eliminate large data centres. It should be possible to have very powerful centralized machines without reducing every other computer to a permission-dependent terminal.

These responses can be unified under the idea of minimum viable computational sovereignty. An independent researcher need not reproduce the capability of a frontier laboratory. What matters is retaining enough independent capacity to formulate hypotheses, run meaningful experiments, inspect results, preserve intellectual work and verify claims made by more powerful institutions. Let $T_i^{\mathrm{core}} \subseteq T_i$ be the tasks necessary to preserve her independent investigative and productive capacity, such as reading and modifying source code, executing experiments, preserving records, reproducing numerical results and inspecting the outputs of external systems. A worker has minimum viable computational sovereignty when

$$\forall t \in T_i^{\mathrm{core}} : \quad \exists S \in \mathscr{S}_i(t) \text{ such that } S \subseteq L_i,$$

that is, when every core task has at least one sufficient configuration under her independent control. The condition does not demand independence in every computational activity. A researcher may depend on a supercomputer for an exceptionally large simulation without depending on it to formulate the model, inspect the equations, run smaller experiments or evaluate the resulting claims. The distinction between frontier capability and independent investigative capability keeps the demand attainable, since

otherwise sovereignty would require every individual to possess the whole technological infrastructure of civilization. Competing directly with the largest infrastructure owners on aggregate expenditure is probably futile, whereas retaining the capacity to investigate, reproduce and challenge their conclusions is a different and more attainable objective. The hobbyist movement did not succeed because every enthusiast owned a machine equal to the largest mainframe. It succeeded because inexpensive machines were programmable enough to support an alternative ecosystem, and the contemporary counterpart would combine smaller models, efficient algorithms, commodity hardware, portable software and cooperatively governed infrastructure.

The forward deployed engineer observation in the video points to an additional channel of concentration, in which services that independent contractors formerly supplied are bundled into the offerings of platform owners (00:34:16 to 00:35:11). Organizational economics explains why such integration might occur, since the boundary between market contracting and firm organization is a function of transaction costs [2], and a provider that holds both the infrastructure and the integration expertise has reasons to internalize both. That analysis does not show that independent contracting is in fact declining as a result, which would require labour market evidence the video does not supply. It is a distinct mechanism from the price of memory, and it would affect independent workers even if component prices were to fall.

10 Conclusion

The video examined here combines a strong structural insight with an overextended claim about intention. The insight is that independence in productive work depends on control over the full set of dependencies that convert labour into income, and that ownership of a single instrument, whether a car or a computer, does not secure it. The overextension is the attribution of deliberate, coordinated dispossession to owners of infrastructure, for which the transcript supplies no evidence and which the structural argument does not need.

The central analytical result of the essay is that an economically efficient allocation of computational resources need not be an allocation that preserves autonomy. Utilization measures how fully installed capacity is used. Autonomy measures how many of a worker's tasks she can begin without permission. Operational availability measures how reliably her attempts are actually served. Dependency measures how much of her work can be disabled by the withdrawal of external resources. Pooling raises the first, may lower the second and may raise the fourth, and the direction of the effect is determined by governance and not by hardware. For this reason apparently redundant, underutilized or decentralized capacity may have considerable social and epistemic value, and an account of the efficient use of scarce components that omits this value is incomplete. Whether the political conclusions that the video draws are warranted is a separate question, which the structural argument leaves open. Three forms of scarcity should be distinguished: physical scarcity, in which resources are insufficient, economic scarcity, in which they exist but are unaffordable, and administrative scarcity, in which they

exist and may be affordable but their use requires authorization that can be withheld. Engineering can relieve physical scarcity without relieving the others, redistribution can relieve economic exclusion without removing administrative dependency, and institutional reform can raise authorization independence without adding a single processor. The personal computer relieved all three at once for a wide class of activities, and a partial reversal need not involve the disappearance of personal computers. It need only involve the migration of consequential capabilities beyond the domain in which individuals can exercise them independently. The organizing observation is that aggregate capability and resource efficiency can rise while independent productive capacity falls, so that technological advance and technological democratization are different things. What the argument does establish is that who holds the authority to initiate computation is a question worth asking before the answer is settled by default.

References

[1] Benkler, Y. (2006). *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. New Haven: Yale University Press.

[2] Coase, R. H. (1937). The Nature of the Firm. *Economica*, 4(16), 386–405.

[3] Dettmers, T., Pagnoni, A., Holtzman, A., and Zettlemoyer, L. (2023). QLoRA: Efficient Finetuning of Quantized Language Models. *Advances in Neural Information Processing Systems*, 36.

[4] Dixit, A. K., and Pindyck, R. S. (1994). *Investment under Uncertainty*. Princeton: Princeton University Press.

[5] Ehrenreich, B., and Ehrenreich, J. (1977). The Professional-Managerial Class. *Radical America*, 11(2), 7–31.

[6] Frantar, E., Ashkboos, S., Hoefler, T., and Alistarh, D. (2023). GPTQ: Accurate Post-Training Quantization for Generative Pre-Trained Transformers. *International Conference on Learning Representations*.

[7] Hinton, G., Vinyals, O., and Dean, J. (2015). Distilling the Knowledge in a Neural Network. arXiv:1503.02531.

[8] Hirschman, A. O. (1970). *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*. Cambridge, MA: Harvard University Press.

[9] Levy, S. (1984). *Hackers: Heroes of the Computer Revolution*. Garden City, NY: Anchor Press/Doubleday.

[10] Liu, C. (2021). *Virtue Hoarders: The Case against the Professional Managerial Class*. Minneapolis: University of Minnesota Press.

[11] Marx, K., and Engels, F. (1848). *Manifesto of the Communist Party*. London.

[12] Marx, K. (1867). *Capital: A Critique of Political Economy*, Volume I. Hamburg: Meissner.

[13] Non-Compete (EJ). *The Hidden Class System Behind the AI Crisis*. Video essay, YouTube. `https://youtu.be/81HjmQjp4_4`. Auto-generated English transcript consulted, with passages located by timestamp.

[14] Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.

[15] Rosenblat, A., and Stark, L. (2016). Algorithmic Labor and Information Asymmetries: A Case Study of Uber's Drivers. *International Journal of Communication*, 10, 3758–3784.

[16] Srnicek, N. (2016). *Platform Capitalism*. Cambridge: Polity Press.

[17] Turner, F. (2006). *From Counterculture to Cyberculture: Stewart Brand, the Whole Earth Network, and the Rise of Digital Utopianism*. Chicago: University of Chicago Press.

[18] Varoufakis, Y. (2023). *Technofeudalism: What Killed Capitalism*. London: Bodley Head.