Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

The Calculation Problem Moves Inside the Firm: AI-Driven Budgeting and the Misesian Limits of Organization

Alfredo Agustin Merlet Echavarria

alfredo.merlet@gmail.com

Universidad del Alba

Research Article

Keywords: economic calculation, artificial intelligence, budgeting, theory of the firm, transfer pricing, knowledge problem

Posted Date: November 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11694656/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

The return of the socialist calculation debate in the age of artificial intelligence has been conducted almost entirely at the level of the whole economy, where Austrian economists have reached a broad consensus that AI does not solve the knowledge problem. This paper moves the debate to the setting in which non-price allocation is the norm: the firm. Building on Coase (1937) and Klein (1996), it treats the budget as the firm's central plan and asks how AI changes the problem of economic calculation where resources are allocated through budgets rather than prices. The paper argues that AI relaxes Simonian bounded rationality but not the Misesian calculation constraint, because forecasting models learn from prices and cannot generate them. Three falsifiable propositions follow, each set against a rival technosocialist prediction: AI's accuracy gains are largest for budget items with an external price reference; AI induces internalization that shows diminishing returns where no external market exists; and models trained only on internal data perpetuate budgetary bias. The analysis yields a counterintuitive result: far from freeing the firm from market prices, AI increases its dependence on them. Illustrations from AI-assisted transfer pricing, electricity markets and Chile's Project Cybersyn are discussed. For managers, the implication is that AI-assisted budgeting should preserve external benchmarks and market-based transfer prices rather than replace them.

JEL classification: B53; D23; M41; O33; P21

1 Introduction

The socialist calculation debate has returned. Advances in machine learning, big data and computing power have revived the claim that a planner equipped with sufficient data and computational capacity could replicate, or replace, the allocative role of market prices (Dapprich and Cockshott 2023). Austrian economists have answered that the original problem identified by Mises ([1920] 1975) and Hayek (1945) was never one of computation. It concerns knowledge that is dispersed, tacit and generated only through market exchange, and the property rights that make such exchange possible (Boettke and Candela 2023; Lambert and Fegley 2023; Phelan and Wenzel 2023).

On this point the Austrian response has reached a broad consensus: artificial intelligence does not solve the knowledge problem at the level of the economy as a whole. Karadimas (2023) argues that economic calculation without markets is epistemically impossible, not merely computationally hard. Gmeiner and Harper (2024) and Martinelli (2025) argue that algorithmic planning cannot reproduce the signals of relative scarcity that market prices convey. Berggren (2026) specifies the conditions under which AI strengthens, rather than weakens, the comparative case for markets. Yet this literature has been conducted almost entirely at the scale of the whole economy. It leaves aside the setting in which non-price allocation is the norm rather than a hypothetical: the firm.

Inside the firm, resources are allocated by budgets rather than by prices (Coase 1937). Klein (1996) showed that the Misesian argument therefore places a limit on the size of organizations: when internal activities lose an external market price against which they can be valued, the firm can no longer calculate

profit and loss for them. The budget is, in this sense, the firm's central plan. It is also a domain to which AI is increasingly applied, from forecasting and budgeting to transfer pricing and management control (Korhonen et al. 2021; Siva 2026). Davidson (2024) extends the institutional analysis of AI to the firm but does not specify the accounting mechanism through which calculation operates. Dabu and Aligica (2026) address strategy and human capabilities rather than resource allocation. The management accounting literature reports that calculation tasks remained better suited to humans than to machines in an actual automation process (Korhonen et al. 2021) and that management accounting retains a maieutic role in non-routine decisions as it goes digital (Tiitola et al. 2024), but it does not draw on the Mises–Klein framework. As a result, the claim that AI faces the same limits inside the firm as in the economy at large has been asserted, but its mechanism remains unexplained.

This paper asks how artificial intelligence changes the problem of economic calculation within the firm, where resources are allocated through budgets rather than prices.

The paper makes three contributions. First, it extends Klein's (1996) analysis of the limits of organization to AI-driven budgeting and treats the budget as the firm's central plan. Second, it derives three falsifiable propositions, each set against the rival prediction implied by the technosocialist position. The first concerns market anchoring: AI's accuracy gains should be largest for budget items that have an external price reference. The second distinguishes Simon from Mises: AI reduces bounded-rationality costs and induces further internalization, but internalization without external markets yields diminishing returns. The third concerns algorithmic calculational chaos: models trained only on internal data perpetuate budgetary bias. Third, the analysis yields a counterintuitive result. Far from freeing the firm from market prices, AI increases its dependence on them. For managers, the implication is that the design of AI-assisted budgeting should preserve external benchmarks and market-based transfer prices rather than replace them.

The paper is conceptual. It develops its argument from the Austrian and transaction-cost literatures and from the management accounting literature on automation, identified through structured searches in Scopus, and it states its propositions in a form that can be tested with budget-line data. The remainder of the paper is organized as follows. Section 2 traces the argument from socialist calculation to calculation inside the firm. Section 3 examines what AI does, and does not do, in budgeting. Section 4 states the three propositions and their falsification conditions. Section 5 offers illustrations from transfer pricing, electricity markets and historical computer-assisted planning. Section 6 discusses implications, a research agenda and limitations, and Section 7 concludes.

2 From socialist calculation to calculation inside the firm

2.1 Mises and Hayek: prices as the condition of calculation

Mises ([1920] 1975) did not argue that a socialist planner would lack the will or the intelligence to allocate resources well. His claim was narrower and more radical. Economic calculation requires a

common unit in which heterogeneous means of production can be compared, and in a monetary economy that unit is the money price of factors. Those prices exist only because factors are privately owned and exchanged. Abolish private ownership of the means of production and the exchange of factors disappears; with it disappear the prices that would have told the planner whether a given combination of steel, labour and land produces more value than it consumes. Consumer goods might still be valued, but the planner could not impute values back to the higher-order goods used to produce them. Without that imputation there is no way to compare the cost of alternative production plans, and therefore no way to know whether capital is being maintained or consumed.

The argument is important for what follows because it locates the problem in the absence of a particular kind of information, not in a shortage of computing power. A planner who could solve any system of equations would still face the problem, because the coefficients of the system are what the market would have revealed and the planner does not have them.

Hayek (1945) added a second layer. The relevant knowledge is not given to anyone in its totality. It is dispersed among millions of individuals as knowledge of the particular circumstances of time and place: an idle machine, a temporary surplus, a supplier who can deliver faster than usual. Much of it is tacit and cannot be transmitted as statistics, and much of it does not exist until people are confronted with a change and decide how to respond. The price system economizes on this knowledge. A rise in the price of a material tells users to economize on it without telling them why, and that signal propagates through the economy faster and further than any central collection of data could. In Hayek's account the market is a discovery procedure and a communication system at the same time.

Bowles, Kirman and Sethi (2017) show that this view anticipated the modern idea of the market as an algorithm: a decentralized procedure that aggregates information which no participant possesses. The metaphor cuts both ways. If the market is an algorithm, it is tempting to conclude that a better algorithm could replace it, which is precisely the technosocialist inference (Dapprich and Cockshott 2023). The Austrian reply is that the market algorithm is not separable from its inputs: the prices it computes are produced by the exchanges of owners who bear the consequences of their decisions. Sunstein (2024) makes a related point about algorithms in society more generally: algorithms improve accuracy by reducing bias and noise in human judgment, but they struggle where the data needed to answer the question do not exist, including where the relevant information is local knowledge held privately by individuals.

Two elements of this tradition matter for the firm. The first is Mises's emphasis on factor prices as the basis for capital accounting. The second is Hayek's emphasis on local, tacit knowledge that only becomes usable through prices. Both reappear, in a different form, inside organizations.

2.2 Coase and Klein: the firm as an island of planning

Coase (1937) asked why firms exist if the price system coordinates economic activity so well. His answer was that using the price mechanism has costs: discovering the relevant prices, negotiating and

enforcing contracts for each transaction. Within the firm, these market transactions are replaced by the direction of an entrepreneur-coordinator. The firm expands until the cost of organizing an additional transaction internally equals the cost of carrying it out through the market or of organizing it in another firm. Inside the firm, in other words, resources move not because a price changed but because someone with authority decided that they should.

This makes the firm a limited version of the socialist planner, and the Austrian tradition has drawn the parallel explicitly. Klein (1996) argued that the calculation problem sets an upper limit on the size of the firm that is distinct from the transaction-cost limits stressed by Coase and his successors. A firm can calculate the profitability of its internal divisions only as long as it can value their inputs and outputs. It does so by reference to external markets: transfer prices based on market prices, opportunity costs measured against what a resource could earn outside, and make-or-buy comparisons against external suppliers. When a firm becomes so integrated that an activity no longer has any external market, it loses that reference. It can still allocate resources to the activity, but it can no longer tell whether doing so creates or destroys value. In Klein's terms, the firm runs into the same calculational chaos that Mises attributed to the socialist commonwealth, only at a smaller scale.

Klein's argument has a precise implication for internal accounting. Transfer prices, internal charges and cost allocations are not neutral bookkeeping devices; they are the firm's way of importing market information across its own boundary. Where they are anchored in external prices, divisional profit is meaningful and the firm can compare uses of capital. Where they are set by cost-plus formulas, negotiation or administrative fiat, the resulting figures are internally consistent but say nothing about opportunity cost. The firm can still produce a complete set of accounts. What it cannot do is know whether those accounts describe value.

2.3 The budget as the firm's central plan

The budget is the instrument through which this internal planning takes place. Its structure is closely analogous to the central plan of the calculation debate, and it is useful to make the analogy explicit because it is the bridge between the Austrian argument and AI.

First, the budget is a plan for the allocation of resources over a defined period, set in advance and from the centre. Headquarters or a budget committee fixes revenue targets, cost envelopes and capital expenditure for each unit. Units then operate within those envelopes, and deviations are measured against the plan rather than against what the resources could have earned elsewhere.

Second, the information on which the budget is built is largely supplied by the units being budgeted. The centre does not know the particular circumstances of each unit; it asks the units to forecast them. This reproduces Hayek's problem in miniature. The knowledge that matters for allocation is local, and the centre must rely on reports that are filtered through the incentives of those who make them.

Third, those incentives are well known to anyone who has taken part in a budgeting cycle. Units that are evaluated against budget targets have reasons to understate revenue and overstate costs so that targets

are easy to meet. Units whose allocation depends on last year's spending have reasons to spend the whole envelope before the period closes. The centre, knowing this, adjusts the submissions, and the result is a negotiation in which the final figures reflect bargaining power and precedent as much as economic information. Incremental budgeting, in which next year's figure is last year's figure plus an adjustment, is the routine form of this process. It is, in effect, a planning procedure that replaces calculation with inertia.

Fourth, internal charges are set within the same process. The cost of shared services such as information technology, finance, human resources or a regional headquarters is distributed to operating units by allocation keys: headcount, revenue, floor space or negotiated fixed fees. When an equivalent service can be bought outside, the internal charge can be compared with an external price and units can contest it. When there is no outside equivalent, the charge is whatever the allocation key produces, and neither the provider nor the user can tell whether the service is worth its cost.

The analogy has limits. Unlike the socialist planner, the firm is embedded in markets: it buys inputs, sells outputs, raises capital and is ultimately tested by profit and loss. This is exactly Klein's point. The firm's internal planning is disciplined to the extent that its budget lines can be referred to external prices, and undisciplined to the extent that they cannot. The question for this paper is what happens to that discipline when the budget is increasingly produced by AI.

3 What AI does, and does not do, in budgeting

3.1 What AI does: forecasting, pattern detection and automation

AI enters budgeting first as a forecasting technology. Machine-learning models trained on historical transactions, sales, prices and macroeconomic indicators can produce revenue and cost forecasts at a level of granularity and frequency that a manual process cannot match. They can update rolling forecasts continuously instead of once a year, detect seasonal and cyclical patterns, flag anomalies in spending, and reconcile large volumes of transactions automatically. Generative tools add the ability to draft variance commentaries and to answer queries about the budget in natural language.

In the terms of the organizational literature, these capabilities relax bounded rationality. Davidson (2024), drawing on Simon's work, argues that AI changes the economic institutions of production because it lowers the cost of processing information and of coordinating decisions. A budget office equipped with AI can consider more variables, more scenarios and more units than one without it, and it can do so faster. The cost of internal coordination falls. In Coasean terms, this should move the boundary of the firm outward: if organizing a transaction internally becomes cheaper, more transactions will be organized internally.

The practice of transfer pricing shows what kind of information these tools actually rely on. Siva (2026) reviews the use of AI in transfer pricing and finds that multinational enterprises apply machine learning mainly to automate benchmarking, documentation and operational transfer-pricing monitoring, while tax authorities use AI-driven risk models for audit selection. In other words, AI is used to search the market for reference prices more efficiently, not to dispense with them. The arm's-length principle requires that internal prices be benchmarked against what independent parties would charge, and AI makes that benchmarking cheaper and more thorough. The value of the tool comes from the external prices it finds.

This pattern generalizes. A forecasting model for the cost of a purchased input is trained on past purchase prices, which are market prices. A model of expected revenue is trained on past sales, which are market transactions. Where the budget line corresponds to something bought or sold outside the firm, the model is learning from prices that already embody the dispersed knowledge Hayek described. Its accuracy is borrowed from the market.

3.2 What AI does not do: prices, tacit knowledge and judgment

The limits of AI in budgeting follow from the same observation. A model can learn the relationship between past prices and other variables; it cannot produce a price where no exchange has taken place. For a budget line that has no external counterpart, such as an internal service with no outside equivalent, the only training data are the firm's own past allocations. These are the output of the budgeting process itself, including its negotiation, its inertia and its allocation keys. A model trained on them learns to reproduce them.

The opportunity cost of a resource illustrates the problem. Opportunity cost is the value of the best alternative use of a resource, and in a market economy it is revealed by what others are willing to pay for it. A model trained on internal data observes how the firm has used the resource, not what it could have earned elsewhere. Unless the training data include external prices, the model has no access to the alternative against which the use should be judged. More data of the same kind do not help, because the missing information is not in the data.

The Austrian literature on AI identifies further limits. Lu (2025) argues that large language models do possess two forms of tacit knowledge, knowledge that is codifiable but too costly to articulate and knowledge of nuance encoded in language, but not embodied knowledge gained through experience, and that they remain subject to the Hayekian constraints of dispersed knowledge and bounded rationality. The concession matters for the argument here: even a model with tacit knowledge of language has no access to the opportunity cost of a resource that has never been exchanged. Kurun (2026) argues that entrepreneurial discovery depends on embodied tacit knowledge that is not available to AI systems, and Martinelli (2026) makes a similar argument about innovation. Zhang (2026) argues that AI does not understand tacit knowledge but bypasses understanding through what he calls functional manifestation, reproducing its effects through probabilistic fitting. Even on this reading, what AI fits is information already expressed in data, not the price signals that the budgeting problem lacks.

Davidson (2025) argues that AI cannot replace the entrepreneurial function because it lacks the judgment and evaluative agency required to act under uncertainty, and McDavid, Kiesling and Chassin (2026) argue that AI agents can participate in markets but depend on market institutions, agency and trust to do so. Wang and Alonso-Neira (2026) reach a parallel conclusion for academic work: AI substitutes for standardized, information-intensive tasks but increases the relative importance of judgment under uncertainty.

For budgeting, the relevant form of judgment is entrepreneurial: deciding that a unit should receive more resources than its history justifies because of an opportunity that is not yet visible in the data, or fewer because of a threat that has not yet materialized. This is the kind of decision that the budgeting process exists to make and that a model trained on the past is least able to make.

3.3 Field evidence: why calculation remained human

The management accounting literature offers direct evidence on these limits, although it does not use the Austrian framework. Korhonen et al. (2021) followed, through an interventionist case study at a machinery manufacturer, an attempt to automate the cost calculations used to price customer tenders. A task that had appeared programmable turned out not to be: the new sales configurator became a layer on top of the old spreadsheet calculation, the cost specialist continued to authorize large deals, and the tool was eventually abandoned. The authors identify the cost specialist's authorization of each tender as the most nonprogrammable element, alongside technical difficulties such as data maintenance and freight cost estimation, and conclude that automating nonprogrammable tasks risks faster but false decisions.

Read through the Mises–Klein lens, this finding admits a specific interpretation, which goes beyond the authors' own framing. The part of management accounting that resists automation is the part that has to supply, by judgment, what the market does not supply by price. When an accountant builds a calculation for a make-or-buy decision or for a new internal service, the hard step is not the arithmetic but the decision about which external reference applies and how to adapt it. Where such a reference exists, the step is easy and can be automated. Where it does not, it requires judgment.

Tiitola et al. (2024) describe the maieutic role of management accounting, by analogy with the Socratic method: in non-routine decisions, accounting helps managers articulate and examine the assumptions behind their choices rather than supplying definitive answers. Analysing practitioners' discourses on AI, they find that this support is fundamentally maieutic, that AI as it currently stands is not expected to change this by providing perfect answers, and that digitalization has the potential to sustain or even enhance the maieutic role. In the terms of this paper, the maieutic role is a substitute for the price system within the firm. Where the market does not discipline an internal decision, the accountant's questioning does: it forces managers to state what an activity is worth and to compare it with alternatives. AI may support that dialogue by producing figures faster, but it cannot conduct it, because the figures it produces for unpriced activities are derived from the same internal history the dialogue is meant to question.

4 Propositions

The argument of Sections 2 and 3 can be stated as three propositions. Each is formulated so that it implies a different observable pattern from the rival view that AI can substitute for prices inside the firm. That rival view is the firm-level analogue of the technosocialist position (Dapprich and Cockshott 2023): if better data and computation can solve the planning problem for an economy, they should a fortiori solve it for a firm, where the planning problem is much smaller.

4.1 Proposition 1: market anchoring

Proposition 1

The improvement in budget forecast accuracy following the adoption of AI is larger for budget lines that have an external market price reference than for budget lines that do not.

Mechanism. Forecasting models learn from their training data. For lines tied to purchased inputs, sold outputs or services with outside equivalents, the training data include market prices that embody dispersed information about scarcity and demand. For lines without an outside equivalent, the training data consist of the firm's own past allocations. The model can improve the extrapolation of the former; for the latter, it can only reproduce more efficiently the process that generated the history.

Rival prediction. If AI substitutes for prices, accuracy gains should be roughly uniform across budget lines, or even larger for internal lines, where manual processes have historically been weakest.

Falsification. The proposition would be refuted by budget-line data, before and after AI adoption, showing that the reduction in forecast error for lines without external references is equal to or greater than that for lines with them, controlling for line volatility and size.

4.2 Proposition 2: Simon yes, Mises no

Proposition 2

AI reduces the cost of internal coordination and therefore induces firms to internalize additional activities, but the performance of internalized activities that have no external market shows diminishing returns relative to those that do.

*Mechanism. AI relaxes bounded rationality (*Davidson 2024*). This lowers the Coasean cost of organizing transactions inside the firm and moves the boundary of the firm outward. But it does not relax the Misesian constraint identified by* Klein (1996)*: activities without an external market reference cannot be valued. As integration proceeds into such activities, the firm gains coordination capacity but loses the ability to calculate their contribution, so their performance deteriorates relative to activities that retain an external reference.*

Rival prediction. On the technosocialist view, the constraint on firm size is computational. AI should therefore allow firms to integrate without a calculational limit, and integrated activities without external markets should perform as well as those with them.

Falsification. The proposition would be refuted if, after AI adoption, firms that integrate vertically into activities without external markets show segment performance equal to or better than firms that integrate into activities with external markets, or if AI adoption is not associated with increased internalization at all. The second outcome would refute the first half of the proposition, the first outcome the second half.

4.3 Proposition 3: algorithmic calculational chaos

Proposition 3

Where AI models are trained only on internal budget data, without external price references, budgetary biases persist or increase after adoption.

Mechanism. Internal budget histories carry the systematic biases of the budgeting process: budgetary slack, incremental inertia and allocation keys unrelated to value. A model trained on these data learns the biases as regularities and reproduces them in its forecasts, now with the authority of an apparently objective method. Because the model's outputs feed the next budget, which then becomes training data, the bias can become self-reinforcing. This is a firm-level analogue of the calculational chaos that Mises predicted for an economy without factor prices: the system produces figures, but the figures do not describe value.

Rival prediction. On the rival view, learning corrects bias: as the model observes outcomes, systematic deviations between budget and actual results should shrink.

Falsification. The proposition would be refuted if measures of budgetary bias, such as the persistence of favourable variances in the same lines or the correlation of allocations with prior-year allocations net of performance, decline after AI adoption in units whose models are trained only on internal data, to the same extent as in units whose models include external references.

Together, the three propositions imply the counterintuitive result stated in the introduction. If they hold, the value of AI in budgeting depends on the presence of external prices, and the firms that gain most from AI are those that maintain the strongest links between their internal accounts and the market.

5 Illustrations

The illustrations in this section are not tests of the propositions. They show that the mechanism proposed is consistent with how AI is being used in practice and with historical experience of computer-assisted planning.

5.1 Transfer pricing: AI as a search for market prices

Transfer pricing is the area of internal accounting in which the link between internal and external prices is most explicit, because tax rules require it. Multinational groups must show that prices charged between affiliates correspond to what independent parties would have agreed. The documentation process consists largely of searching for comparable transactions and companies, adjusting them for differences, and establishing an arm's-length range.

Siva (2026) shows that AI is being adopted in this field to perform exactly this search, with measurable gains in faster benchmarking and documentation and in predictive compliance controls. The risks he identifies concern opaque 'black box' outputs, data and algorithmic bias, unsettled regulatory treatment and disclosure disputes, and he concludes that AI is a net opportunity only when deployed as decision support under human-in-the-loop safeguards. What AI does not do in this setting is replace the arm's-length benchmark with an internally computed price. Its function is to bring more market information into the firm, more cheaply.

This is the pattern predicted by Proposition 1. Where a market reference exists, AI makes it easier to use. The illustration also points to the boundary of the argument. For intra-group transactions in intangibles or highly integrated services, comparables are scarce, and transfer pricing falls back on profit-split and cost-based methods. These are precisely the cases in which the internal price becomes an administrative construct, and in which the proposition predicts that AI will add least.

5.2 Electricity markets: machine learning within a market

Electricity markets offer a second illustration, at the boundary between firm and market. They combine intensive use of algorithmic forecasting with price formation through exchange. Oğuz and Peker (2025) ask whether machine learning can reduce volatility in electricity markets. They argue that machine learning helps by producing more accurate price forecasts that guide market participants, but that it cannot replace the coordinating role of market processes and price signals, and that it remains weak in forecasting customer behaviour in decentralized, renewable-driven systems.

The relevance for budgeting is that even in an industry where data are abundant, measurement is continuous and the physical system is highly engineered, machine learning has been used to improve decisions taken in response to prices, not to eliminate prices. If AI does not displace price formation in a market where the conditions for algorithmic planning are unusually favourable, there is little reason to expect it to do so inside a firm whose internal activities are far more heterogeneous.

5.3 Project Cybersyn: a historical antecedent

Chile's Project Cybersyn, developed under the Allende government in the early 1970s, is the best-known historical attempt to use computers and real-time information flows to coordinate production in nationalized enterprises. Lehto and Paniagua (2025) study it alongside the subsequent neoliberal reforms as twin experiments in institutional design. Its statistical software, Cyberstride, processed data

on firms' inputs and outputs transmitted through a telex network and sent a warning signal to higher levels of government whenever a variable fell outside the default values set by the system (Lehto and Paniagua 2025). Nothing in this architecture formed prices for the resources it monitored.

Cybersyn is an antecedent, not evidence. Its brief life and the political circumstances of its end make it impossible to draw empirical conclusions about the feasibility of computerized planning from it. It is relevant here for a narrower reason: its design anticipated the architecture of contemporary AI-assisted budgeting. It monitored quantities and variances against a plan, and it relied on managers to interpret deviations. Like an internal budget model, it could tell its users that something had diverged from the plan but not whether the plan itself allocated resources to their most valuable uses. That question requires a reference outside the plan.

5.4 What the illustrations show

The three cases share a pattern. Where AI or computer-assisted coordination operates alongside market prices, as in transfer pricing and electricity markets, it adds value by making price information easier to obtain and use. Where it operates without them, as in Cybersyn and in the parts of transfer pricing that lack comparables, it can monitor and coordinate but cannot value. This is the distinction between relaxing the Simonian constraint and relaxing the Misesian one. The illustrations do not establish that the propositions hold, but they make clear what kind of evidence would bear on them.

6 Implications and research agenda

6.1 Implications for managers

Four implications follow for managers and chief financial officers who are introducing AI into budgeting.

First, the design of AI-assisted budgets should preserve and, where possible, strengthen external references. Budget lines should be classified according to whether they have a market benchmark, and models for lines that have one should be trained on external prices and not only on internal history. Investment in external data, such as supplier price indices, market rates for services and comparables for shared services, is likely to raise the return on AI more than additional modelling capacity.

Second, internal charges should be anchored to market prices wherever an outside equivalent exists. Market-based transfer prices and periodic make-or-buy reviews do more than allocate costs: they import information about opportunity cost into the firm. AI lowers the cost of obtaining such benchmarks and should be used to extend them, not to replace them with model-generated allocations.

Third, model outputs for unpriced activities should be treated as descriptions of the firm's past allocation, not as estimates of value. Where a budget line has no external reference, the model forecast tells managers what the firm would spend if it continued to behave as before. Decisions about whether to expand, shrink or outsource such activities still require judgment and, where possible, the construction of an external reference through market testing.

Fourth, the maieutic role of management accountants should be protected as routine work is automated (Tiitola et al. 2024). The time released by automation is most valuable if it is redirected to questioning the assumptions behind unpriced allocations, which is the part of calculation that remains human (Korhonen et al. 2021).

6.2 Implications for policy

The argument has a modest policy implication. Debates on AI regulation often assume that AI will either concentrate economic power in large integrated organizations or make central coordination more effective. If the propositions hold, the efficiency gains from AI in organizations depend on the continued existence of external markets against which internal activities can be valued. Policy that preserves competitive markets for intermediate goods and services, and accounting and tax standards that require market-based benchmarks, such as the arm's-length principle, therefore support the productive use of AI rather than constrain it. This is consistent with arguments that the governance of AI should rely on decentralized, polycentric arrangements that preserve the use of local knowledge (Makridis and Lazanski 2025) and that blanket precautionary approaches risk suppressing the discovery processes through which the value of new technologies is learned (Phelan 2025).

6.3 Research agenda

The propositions call for data at the level of the budget line, which are rarely available in public sources. Testing Proposition 1 requires panels of budgeted and actual figures by line, before and after the adoption of AI forecasting, with each line classified by the presence of an external price reference. Such data could be obtained through collaboration with firms or with providers of planning software, or from public-sector budgets, where line-level data are more often published and where adoption of AI forecasting can be dated. Proposition 2 can be approached with segment reporting and measures of vertical integration in listed firms, combined with indicators of AI adoption from disclosures or job postings, and with a classification of segments by the existence of external markets for their outputs. Proposition 3 requires measures of budgetary slack and incremental inertia over time and information on the training data of the models used, which points to case studies and surveys of finance functions as the most practical first step. Comparative work across countries would also be valuable, since firms in emerging economies often operate with thinner external markets for intermediate services, which on this argument should limit the benefits they obtain from AI in budgeting.

6.4 Limitations

The paper has three main limitations. First, it is conceptual: the propositions are derived from theory and supported by illustrations, not tested. Their value depends on future empirical work, and some of the data required may be difficult to obtain. Second, the classification of budget lines into those with and without an external market reference is a simplification. In practice, references vary in quality, and many lines are partially anchored, for example through benchmarks for a component of the cost. A full test would need a graded measure. Third, the argument treats AI as a forecasting and coordination technology trained on historical data. Future systems that act as agents in external markets, as McDavid

et al. (2026) discuss, could create new sources of price information within the firm. That possibility would not overturn the argument, since such agents would draw on market exchange, but it could change the boundary between priced and unpriced activities.

7 Conclusion

The return of the calculation debate has focused on whether AI can replace markets in the economy as a whole. This paper has argued that the more immediate question concerns the firm, which already allocates resources through a central plan in the form of its budget. AI makes that plan cheaper to produce and more detailed, relaxing the bounded rationality that limits internal coordination. It does not supply the external prices on which economic calculation depends, because forecasting models learn from prices and cannot generate them.

The consequence is a reversal of the technosocialist expectation. AI does not free the firm from the market; it increases the value of the market information the firm can import across its boundary. The firms that gain most from AI in budgeting should be those that keep their internal accounts most closely anchored to external prices, while those that use AI to extend planning into activities without external markets should find that their models reproduce, with greater confidence, the biases of their own past allocations. Mises's argument does not disappear inside the firm. It becomes a design principle for how AI should be used there.

Declarations

Funding.

No funding was received for conducting this study.

Competing interests. The author has no competing interests to declare that are relevant to the content of this article.

Use of AI tools. Claude (Anthropic; Claude Opus 5.5, accessed November 2026) was used to support the screening of Scopus records, to check bibliographic details of references against primary sources, to draft text in English from the author's argument map and notes, and for language editing. The author designed the argument, selected the sources, reviewed and revised all text, and takes full responsibility for the content of the manuscript. No AI tool is listed as an author.

Data availability. Not applicable. No datasets were generated or analysed in this study.

Author Contribution

A.A.M.E. conceived the study, conducted the literature search, developed the argument and propositions, wrote the manuscript, and read and approved the final version.

References

1. Berggren, N. (2026). When does artificial intelligence strengthen the case for markets? Revisiting the socialist calculation debate. *Public Choice*. https://doi.org/10.1007/s11127-026-01487-2
2. Boettke, P. J., & Candela, R. A. (2023). On the feasibility of technosocialism. *Journal of Economic Behavior & Organization*, *205*, 44–54. https://doi.org/10.1016/j.jebo.2022.10.046
3. Bowles, S., Kirman, A., & Sethi, R. (2017). Retrospectives: Friedrich Hayek and the market algorithm. *Journal of Economic Perspectives*, *31*(3), 215–230. https://doi.org/10.1257/jep.31.3.215
4. Coase, R. H. (1937). The nature of the firm. *Economica*, *4*(16), 386–405. https://doi.org/10.1111/j.1468-0335.1937.tb00002.x
5. Dabu, A., & Aligica, P. D. (2026). Human-centered strategy in the age of intelligent machines: Reassessing organizational capabilities, human agency, and firm-level entrepreneurship in the era of AI. *The Review of Austrian Economics*. https://doi.org/10.1007/s11138-026-00738-1
6. Dapprich, J. P., & Cockshott, W. P. (2023). Input-output planning and information. *Journal of Economic Behavior & Organization*, *205*, 412–422. https://doi.org/10.1016/j.jebo.2022.10.043
7. Davidson, S. (2024). The economic institutions of artificial intelligence. *Journal of Institutional Economics*, *20*, e20. https://doi.org/10.1017/S1744137423000395
8. Davidson, S. (2025). The limits of artificial intelligence. *The Review of Austrian Economics*. https://doi.org/10.1007/s11138-025-00705-2
9. Gmeiner, R., & Harper, M. (2024). Artificial intelligence and economic planning. *AI & Society*, *39*(3), 985–1007. https://doi.org/10.1007/s00146-022-01523-x
10. Hayek, F. A. (1945). The use of knowledge in society. *The American Economic Review*, *35*(4), 519–530.
11. Karadimas, P. (2023). The epistemic impossibility of economic calculation. *Synthese*, *202*(6), 172. https://doi.org/10.1007/s11229-023-04381-2
12. Klein, P. G. (1996). Economic calculation and the limits of organization. *The Review of Austrian Economics*, *9*(2), 3–28.
13. Korhonen, T., Selos, E., Laine, T., & Suomala, P. (2021). Exploring the programmability of management accounting work for increasing automation: an interventionist case study. *Accounting Auditing & Accountability Journal*, *34*(2), 253–280. https://doi.org/10.1108/AAAJ-12-2016-2809
14. Kurun, I. (2026). Artificial intelligence, entrepreneurial discovery, and embodied tacit knowledge. Philosophy & Technology. https://doi.org/10.1007/s13347-026-01159-5
15. Lambert, K. J., & Fegley, T. (2023). Economic Calculation in Light of Advances in Big Data and Artificial Intelligence. *Journal of Economic Behavior & Organization*, *206*, 243–250. https://doi.org/10.1016/j.jebo.2022.12.009
16. Lehto, O., & Paniagua, P. (2025). Designing freedom: Allende, Pinochet and the twin experiments in cyber-socialism and neoliberalism. *Economy and Society*, *54*(2), 334–358. https://doi.org/10.1080/03085147.2025.2513800

17. Lu, J. (2025). Tacit knowledge in large language models. *The Review of Austrian Economics*. https://doi.org/10.1007/s11138-025-00710-5
18. Makridis, C., & Lazanski, D. (2025). Navigating the knowledge problem in the era of open AI and polycentric governance. *The Review of Austrian Economics*. https://doi.org/10.1007/s11138-025-00708-z
19. Martinelli, E. (2025). An agent-based approach to the limits of economic planning. *AI & Society*, *40*(5), 3391–3403. https://doi.org/10.1007/s00146-024-02078-9
20. Martinelli, E. (2026). Innovation Is by Humans for Humans. Commentary on 'Artificial Intelligence, Entrepreneurial Discovery, and Embodied Tacit Knowledge' by Ismail Kurun. *Philosophy & Technology*, *39*(4), 206. https://doi.org/10.1007/s13347-026-01226-x
21. McDavid, B., Kiesling, L., & Chassin, D. (2026). Markets, agency, and trust: AI agents and the knowledge problem. *The Review of Austrian Economics*. https://doi.org/10.1007/s11138-025-00711-4
22. von Mises, L. (1975). Economic calculation in the socialist commonwealth (trans: Adler S). In F. A. Hayek (Ed.), *Collectivist economic planning: critical studies on the possibilities of socialism* (pp. 87–130). Augustus M. Kelley. (Original work published 1920).
23. Oğuz, F., & Peker, M. Ç. (2025). Can machine learning reduce volatility in electricity markets? Lessons from the economic calculation debate. *Economic Affairs*, *45*(1), 62–77. https://doi.org/10.1111/ecaf.12686
24. Phelan, S. E. (2025). Alternatives to the precautionary principle: An Austrian approach to regulating generative artificial intelligence. *The Review of Austrian Economics*. https://doi.org/10.1007/s11138-025-00676-4
25. Phelan, S. E., & Wenzel, N. G. (2023). Big Data, Quantum Computing, and the Economic Calculation Debate: Will Roasted Cyberpigeons Fly into the Mouths of Comrades? *Journal of Economic Behavior & Organization*, *206*, 172–181. https://doi.org/10.1016/j.jebo.2022.10.018
26. Siva, P. (2026). AI in Transfer Pricing: Risk or Opportunity? *Intertax*, *54*(6–7), 445–457. https://doi.org/10.54648/taxi2026050
27. Sunstein, C. R. (2024). The use of algorithms in society. *The Review of Austrian Economics*, *37*(4), 399–420. https://doi.org/10.1007/s11138-023-00625-z
28. Tiitola, V., Jalonen, T., Rantanen-Flores, M., Korhonen, T., Ruusuvuori, J., & Laine, T. (2024). Discourse analysis on sustaining the maieutic role when management accounting goes digital. *Qualitative Research in Accounting & Management*, *21*(2), 140–164. https://doi.org/10.1108/QRAM-11-2022-0198
29. Wang, W. H., & Alonso-Neira, M. Á. (2026). Artificial intelligence, the knowledge problem, and academic judgment: an Austrian perspective. *The Review of Austrian Economics*. https://doi.org/10.1007/s11138-026-00747-0
30. Zhang, X. (2026). Beyond Hayek: Computing the incomputable—tacit knowledge and power order in the age of digital intelligence. *Journal of Innovation & Knowledge*, *17*, 101050.

https://doi.org/10.1016/j.jik.2026.101050