AI DISCLOSURES PROJECT
A Project of Code for Science & Society

ESSAY · ASIMOV'S ADDENDUM

Where Does Intelligence Reside?

Why AI needs markets — and how specialized production / open protocols can enable this.

Ilan Strauss

July 10, 2026

> This is the full version of a five-minute "lightning" talk given at O'Reilly Media's Foo Camp on Saturday June 27, 2026 at Lighthaven, Berkeley — co-sponsored by the AI Disclosures Project.

PUBLISHED IN Asimov's Addendum — a publication by Tim O'Reilly and Ilan Strauss
CANONICAL asimovaddendum.substack.com/p/where-does-intelligence-reside
DOI https://doi.org/10.59350/r17cg-t7t11 (concept)
VERSION DOI https://doi.org/10.59350/mf5t0-aq536 (v1)
ARCHIVED Rogue Scholar · record jpvfn-fjq66
LICENSE Creative Commons Attribution 4.0 International (CC BY 4.0)

CITE AS

Strauss, Ilan. "Where Does Intelligence Reside?" *Asimov's Addendum*, July 10, 2026. Talk delivered at O'Reilly Foo Camp, Berkeley, June 27, 2026. https://doi.org/10.59350/r17cg-t7t11.

Where Does Intelligence Reside?

Why AI needs markets — and how specialized production / open protocols can enable this.

Ilan Strauss · July 10, 2026 · Asimov's Addendum

> *This is the full version of a five-minute "lightning" talk given at O'Reilly Media's Foo Camp on Saturday June 27, 2026 at Lighthaven, Berkeley — co-sponsored by the AI Disclosures Project.*

1. Introduction

Will or won't AI lead to a "jobless apocalypse"? I want to offer one framework for thinking through this question based on Adam Smith and Allyn A. Young — and Hal Varian on standards.

The question I want to ask you today is as follows: Where does intelligence reside? Where does it live?

The argument I want to make is that the artificial intelligence found in large language models (LLMs) is inert and centralized in the models' weights. This is a very different type of intelligence to that exhibited by markets, which is fundamentally decentralized and constantly being updated.[1] Markets already extend and enable machine intelligence; but often sporadically and not always via market exchange. Integrating these two types of intelligence, including through the price mechanism, is vital if AI is to advance specialized production, market exchange, and ultimately employment creation.

2. Is this Market Intelligence?

To understand market forms of intelligence I want to begin by asking whether this picture below reflects a market-based intelligence. It shows women making men's suits using sewing machines. It looks organized. But according to what principles?



Figure 1. Suits for men are manufactured at the Bolshevichka garment factory (1967). Source: RIA Novosti archive, image #901609, via Wikimedia Commons (CC BY-SA 3.0).

How do these workers know how many suits to make, what types to make, and with what machinery? By informed guesses of a central committee, it turns out.

The above picture is from the Bolshevichka garment factory in Moscow, taken in 1967 when it was part of the Soviet Union. This factory used to produce the majority of men's suits for Russia. Although the Soviet Union excelled in heavy industry production, it ultimately failed in consumer goods production, where market signals were needed to inform changing tastes, quantities, efficiencies in resource utilization, and fundamental innovations. Waste was abundant. Human need could not be met. Production techniques were outdated.[2]

Bolshevichka still exists as a menswear brand and company today, selling men's suits, but now informed by market signals.

2.1 Two core features of market-based intelligence

What would it mean if the above factory was instead organized on the basis of a market-based intelligence? This would require two things.

The first is **decentralized and permissionless innovation**. Markets are a federated system, where information remains dispersed along with the trillions of independent buying and selling actions that it enables. Market-processed information is what enables and disciplines these actions (along with society's conventions and laws).

The second is that **current information signals inform market coordination through constantly shifting agents' incentives**. Prices and product information are always being replenished as preferences and technology change, and demand and supply decisions respond accordingly. Market signals inform decision making but also respond to it. For this to work, price signals and other relevant information must be widely available (and not gameable).

It appears to me that digital markets have till now often tried to mimic market forms of intelligence. Google Search of old is an example of a market-like intelligence, I would propose. Google Search has:

- **Permissionless innovation** (building on the web's open protocols). Anyone can create a website. Information (websites) remain federated on servers.
- **Dynamically updated signals**. What the user clicks on (consumes) in the search result page is used to improve and constantly update website search result rankings. This in turn improves the next allocations of user attention to information, which in turn informs what information websites and other producers should supply to the market.

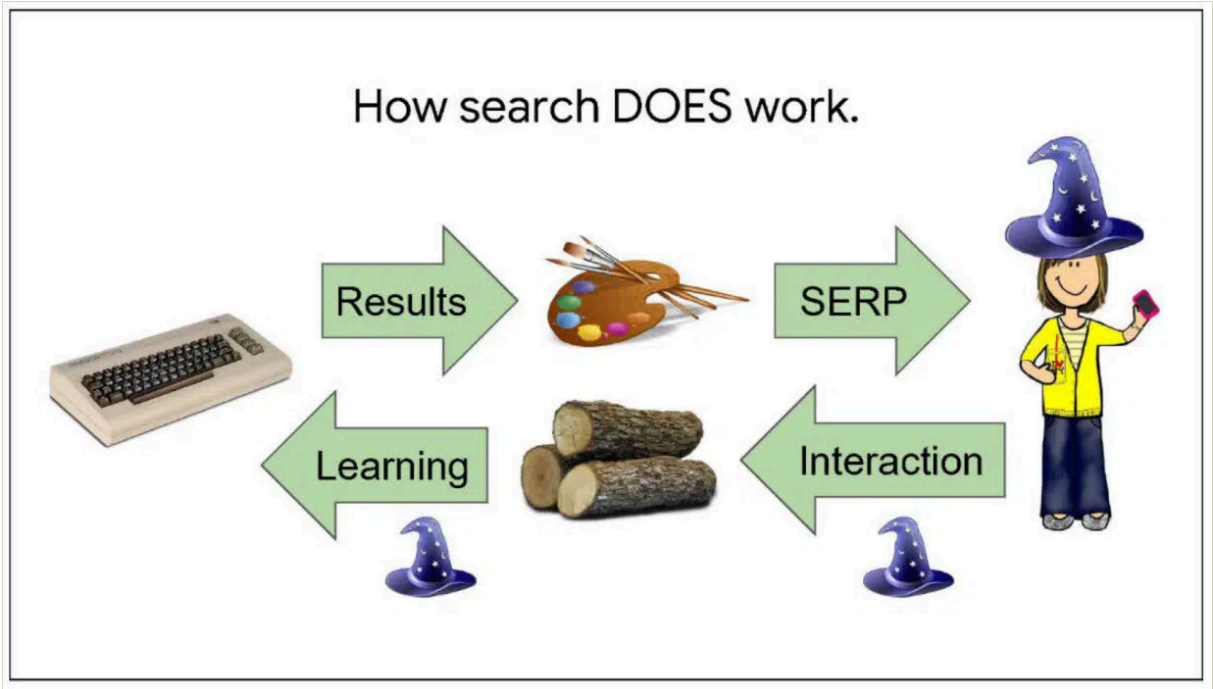


Figure 2. Google internal slide on how search actually works: a feedback loop of SERP → Interaction → Learning → Results. Source: Search Engine Land.

The algorithms of this era advanced what Tim O'Reilly calls "an architecture of participation".

But will markets create jobs? The truth about market-based systems of intelligence is that humans may not always be needed. Nissan's "intelligent factory" below is just one example of labor

being made redundant in manufacturing production. But without markets expanding production on a progressive and balanced basis, little prospect for human employment and prosperity exist.



Figure 3. Robotized automotive assembly — illustrative of Nissan's "intelligent factory" concept. See original video footage from Fox 5 New York (2021).

3. Artificial Intelligence

How do market forms of intelligence compare with present day artificial intelligence in LLMs?

1) An LLM's intelligence is inert and unresponsive at its core. The model's weights are trained on information and feedback after which they become fixed, "dead" things. They are unable to dynamically update or learn as they go. That is why LLMs are reliant on markets for a steady stream of new information, and in particular markets online.

Besides for the internet, the other primary source of live information for models is human users and firms. Every time a user has a conversation with a generative AI system, uploads a document, or tells it something, they allow the model's weights to respond to live, up-to-date, information. Post-training also incorporates human judgment and information, organized by firms such as Surge AI. But direct market signals seem to only weakly inform these activities and the piece rates offered to human annotators.



Figure 4. Human annotators at a post-training data vendor. Post-training internalizes market-like signals within a handful of firms and farms them out via weakly informed piece rates.

2) The other aspect of present AI is that it aims to be generalized — as in artificial general intelligence. In doing so, LLM companies strive to be able to do every task that a human can do and at a high level of competency. This potentially undermines the idea of specialized producers existing across a market, each using different and specialized production techniques. The idea of a single generalized producer ("AGI") is so radical that it threatens to undermine the notion of markets altogether.

And that is the central contention of this talk: that **specialization underpins the ability for a market's intelligence to grow, for permissionless innovation to flourish, and for price signals to function effectively.**

Specialized production provides a bridge between market forms of intelligence and artificial intelligence. So if we want to grow the market for AI, as a system of permissionless innovation, with room for human contributions, it needs to permit — and admit — that specialized production is necessary and desirable. General intelligence divorced from market forms of intelligence cannot survive, except as a stale set of encyclopedias sitting on the shelf.

4. Specialization

Specialized production — or specialization for short — is the ultimate basis for market forms of intelligence and its expansion (an idea we owe to Adam Smith). Specialization is what drives

production for exchange, as for-profit production, coordinated via the market.

Allyn Young in 1928 went one step further when he argued that specialization in the market is a *self-reinforcing* dynamic. The more I specialize as a firm, the more I allow other firms to specialize too, as they can then meet more of their needs (procure more of their inputs) via the market, instead of through internal production. This means that specialization is not only the foundation of market exchange (Adam Smith), but can create an ever-expanding basis for the market itself to grow through increases in output (supply).

Specialization as external economies. The benefits of specialization are best viewed from the industry or ecosystem as a whole that it helps to expand. The early American auto industry is a good example of this. Ford and its rivals began as assemblers, drawing on Detroit's dense ecosystem of specialized machine shops and parts makers. The Dodge Brothers famously supplied Ford's engines and transmissions before ever building a car of their own. It was this external supplier base producing standardized components for a variety of buyers that allowed for mass auto production to arise and prices fall. This grew the size of the market for everyone. Ford later advanced vertically integrated production in its famous River Rouge complex (which Charlie Chaplin's 1936 film *Modern Times* drew on), until more distributed supplier networks gained favor again. The benefits of specialization, therefore, are about the unfolding division of labor between firms, not about a single firm scaling up its own production, which can potentially be to the exclusion of others.

Specialization and integrated / monolithic producers. Specialization undermines the integrated producer, where a single firm attempts to do everything under its own roof, without procuring standardized inputs via the market (or transferring its intermediate outputs only internally, rather than selling them).[3]

Advancing specialization online and in LLMs' own production process can have similar benefits. But this can only have a chance of succeeding if model building is no longer undertaken from start to finish in relative secrecy (housed within a single firm). And if information and intelligence can permanently reside *outside* of a model's weights. Model weights don't only have to consume markets, they can also enable them. Standards and open protocols, as the technical analogs of market-enabled specialized production, are two important mechanisms for advancing such an agenda.

5. Integrating AI with Markets: An architecture of participation for AI

Standards and open protocols are two key mechanisms for advancing market forms of intelligence online that are decentralized and permissionless. The goal is to allow for artificial intelli-

gence to be extended into the living market of information, admitting formal prices and for market-based allocations.

5.1 Standards

Standards enable specialization by allowing interconnected parts to fit together. In markets with many complementary parts, standards are vital to allowing a single producer to address the whole differentiated market of firms as if it were a single source of demand. This helps the supplier to not be captured by a single producer ("lock-in"). As my colleague Tim O'Reilly remarked to me, one much-underappreciated ingredient to the industrial revolution's take-off was specialized inputs and parts being produced to a standard.

Such standards are especially important in networked markets since connecting to the network is vital to tap into its scale and capabilities. This means that it can potentially let average revenue per user grow for any firm with access to the network's expanding utility.

So in technology markets, standards not only facilitate lower costs through specialization / scale, they also inform whether the value of a network is captured by one firm or shared by many.

5.2 Open protocols

Open protocols are shared rules that allow independent users, agents, or devices to coordinate and communicate on the basis of a publicly available specification. They are a type of standard.

By hardening an interface at a given layer through an open specification they allow other firms to connect to it reliably and then specialize against it. They are a battering ram against the integrated producer and the monolithic architecture. A single, closed architecture is designed so that potentially separable complements cannot emerge as independent markets. Under a monolithic architecture specialized production is impeded because the interfaces between those functions simply don't exist or are controlled by a single firm. But a modular architecture is the technical enabler of specialized production that can be disassembled safely (an idea that traces back to Herbert Simon's 1962 paper, "The Architecture of Complexity").

Once specialization takes hold at one interface via an open protocol, the integrated architecture of the AI lab or the platform can be challenged via the market as firms can now compete to supply or draw on discrete components (such as memory, retrieval, payments, identity, tool invocation, attribution, safety filters, and so on).

In doing so, an open protocol can drive vertical disintegration. Interfaces become standardized and what used to be one firm's internal operation gets spun out into a specialized industry as the protocol now handles the coordination (or the identification, the encryption, the security)

that previously required all functions to be retained inside the same firm. That spinning-out of tasks (in economic terms) is precisely Allyn Young's progressive subdivision of production across specialized firms.

Although not usually described in market-shaping terms, open protocols, therefore, are exactly that. They shape the market intelligence that can be built on top of them.

The Model Context Protocol (MCP) comes to mind here. It helped catalyze specialized production in AI applications and a more permissionless innovation by allowing AI agents (clients) to connect to external tools and context. A large ecosystem of specialized servers, each exposing one or more tools and resources, grew up around it, which any compliant client can use. Such specialization in turn grew the market for other specialized producers. Specialization begot specialization and rival labs eventually adopted the standard.

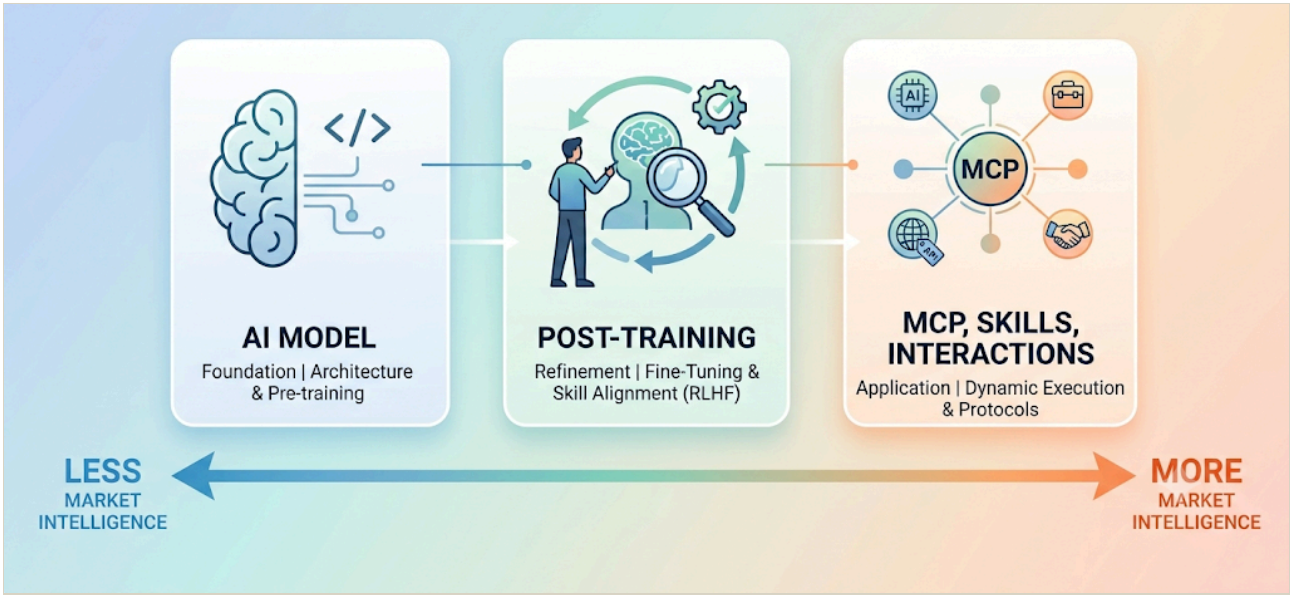


Figure 5. Can intelligence ever exist in isolation? Post-training reflects markets being internalized within a handful of firms and farmed out to workers via weakly informed market-determined piece rates. MCP enables more permissionless, decentralized action and innovation, with the potential for responsive signals / telemetry. Skills are an extension of human intelligence that is valuable when not captured by the model's weights.

6. Concluding Initiatives

The dream of artificial intelligence being very "general" in nature (as in "AGI") seems, at the end of the day, to be technically compatible with a specialized division of labor for its production and with the existence of specialized AI producers. The market for AI startup financing seems to think so too, funding different approaches to AGI (and recursive self-improvement) and funding differentiated, highly specialized AI producers at the same time.

Extending machine intelligence into markets, and vice-versa, to allow for human participation and reward via markets ("an architecture of participation" for AI) is the focus of the three initiatives we are launching this year at the AI Disclosures Project:

- **Open Memory Protocol Initiative (OMPI).** A multi-stakeholder nonprofit consortium developing open, interoperable specifications for portable AI memory. A vendor-neutral open-source ecosystem, with open working groups, it develops an open agentic memory protocol, focused on transferring a core memory object, and allowing for it to be read and written to concurrently (shared memory). Read more: here.
- **The AI Markets Initiative.** A cross-institution economics research and data collaboration that galvanizes the best economic minds around AI market design challenges and opportunities. The goal is to create functioning AI markets with architectures of participation.
- **Skills Marketplace.** Launching a skills marketplace with monetization for skills authors based on modular, interoperable, rights-protected skills, with proven uplift to agentic performance. All specifications and code will be open-sourced, to be used as a prototype for others to adopt.

If you are interested in being involved in any of these three initiatives in whatever capacity, we would love to hear from you. Thank you for your support.

Postscript: Specialization via mixture of experts (MoE) models and agent orchestration are not dealt with here but are worth noting.

Endnotes

1. This does not preclude the role of firms as owner-producers, hierarchical organizers of capitalist production but still ultimately beholden to the market. ↩
2. Notes Richard Ericson, in "The Classical Soviet-Type Economy: Nature of the System and Implications for Reform," *Journal of Economic Perspectives* 5(4), 1991, pg. 21: "When the system pursues a few priority objectives, regardless of sacrifices or losses in lower priority areas, those ultimately responsible cannot know whether the success was worth achieving. The central authorities lack the information and physical capability to monitor all important costs — in particular, true opportunity costs — yet they are the only ones, given the logic of the system, with a true interest in knowing such costs." And then concluding, pg. 26: "Finally, it must be remembered that the ultimate configuration of institutions and interactions is unknowable, a largely unintended consequence of the growth of decentralized agent interaction. Thus, a final lesson for successful reform taught by the nature of the traditional Soviet-style system is to abandon the Faustian urge to control, to know in advance, and thus to allow economic outcomes to arise naturally as the unpredictable consequences of market interaction." ↩
3. So-called agentic orchestration acting on the basis of user memory and context is one attempt to recreate this division of labor arguably, with the user or firm's information providing the specialized ingredients into the

produced service — services are produced as they are consumed. But another approach is simply specialized agents, which implies a more specialized production. The idea that agents lower transaction costs and so might naturally see more production internalized within the firm seems highly speculative. Coordination within a firm is often less efficient / more costly than via the market, noted by Stigler in his 1951 piece ("The Division of Labor is Limited by the Extent of the Market"). The boundaries of specialized production are likely in these instances to be decided more by the sensitivity and accessibility of data — and intellectual property more generally — along with the ability to have standardized parts and a large enough addressable market. ↩

Asimov's Addendum is a publication by Tim O'Reilly and Ilan Strauss on participatory, decentralized AI markets built for human participation, reward, and oversight. This essay is co-published with the **AI Disclosures Project**, a project of Code for Science & Society.