Economics for the Coasean Singularity

Andrey Fradkin (Boston University and Amazon)

This presentation includes joint work with **Rohit Krishnan**, Peyman Shahidi, Gili Rusak, Benjamin Manning, John Horton, Mert Demirer, and Nadav Tadelis.

The views expressed in this talk do not represent the views of Amazon.

Let's seriously consider the idea that AI agents will participate in the economy like human agents.

The Coasean singularity (strawman)

AI agents cause transaction costs to go to zero, resulting in a capitalist utopia and the dissolution of firms (Arrow-Debreu markets).

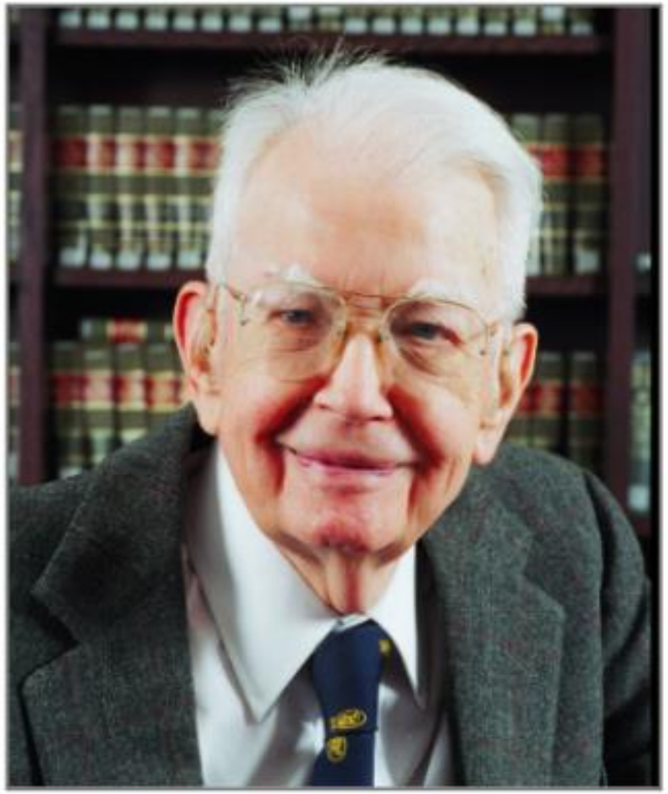

The Nature of the Firm

By R. H. Coase

Economic theory has suffered in the past from a failure to state clearly its assumptions. Economists in building up a theory have often omitted to examine the foundations on which it was erected. This examination is, however, essential not only to prevent the misunderstanding and needless controversy which arise from a lack of knowledge of the assumptions on which a theory is based, but also because of the extreme importance for economics of good judgment in choosing between rival sets of assumptions. For instance, it is suggested that the use of the word "firm"

Statements about AI agents in the economy

1. They will participate as laborers doing productive work.

2. They will act as personal agents engaging in the market on our behalf.

3. They will collaborate, coordinate, and compete in a variety of ways, just like humans do.

4. Coase's insights hold; transaction costs will dictate the organizational structure for both humans and agents.

5. We will need rules, tools, and market design innovations to make this new world work well.

Roadmap

1. **How existing scaffolding for AI agents works.**
2. Why a market-based scaffolding is appealing.
3. Experiments with market-based scaffolding.
4. Market design and concluding thoughts.

What is a scaffolding?

- Software architecture around a language model.

- Determines what happens to the context, and the rules for using tools and calling sub-processes.

- You can think of it as an organizational structure.

Long-running Codex Agent

Big efficiency gains from agents being able to chain tasks

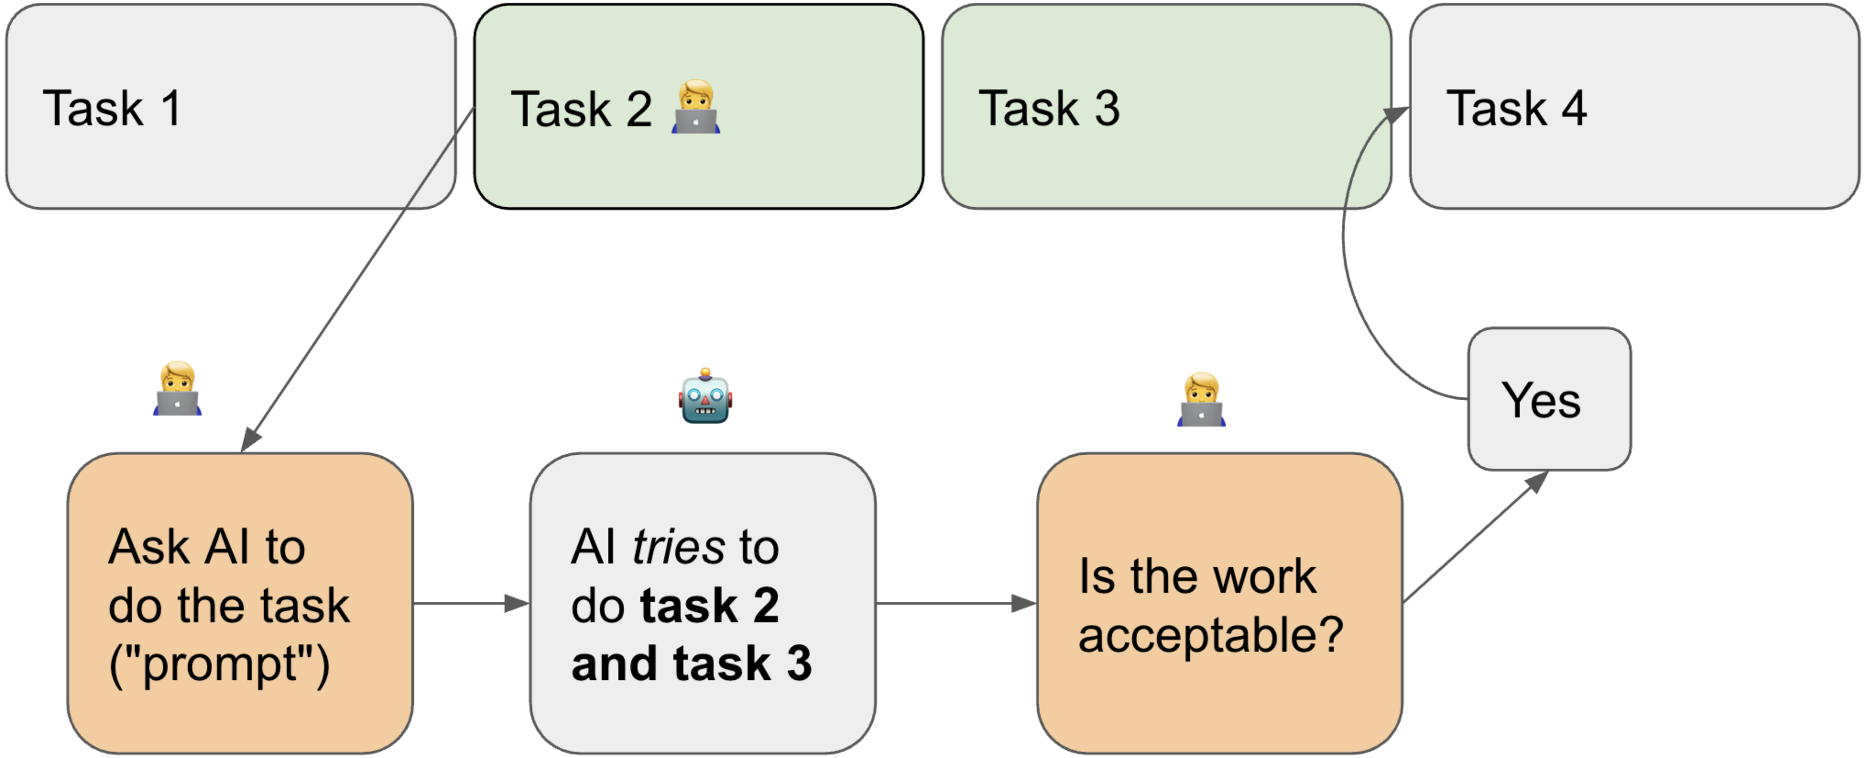

Example: Claude code

- Markdown files specify persistent memory and usage patterns.

- The main instance determines when to spawn sub-agents (for context window protection, parallelization, task isolation, and specialization).

- Sub-agents have their own context and often use a cheaper model (Haiku).

- The process is a while loop that calls sub-agents of depth one.

Other patterns

- Hierarchical sets of instructions (Codex).
- Explicit roles in team-based collaboration (Codex).
- Multiple isolated and persistent agents run in parallel (OpenClaw).
- Cost-aware model selection using deterministic rules (OpenClaw).
- OpenClaw is the only one intended to be used with models from a variety of creators.

Parallels to economic concepts

- The finite context window approximates limited attention.
  - Communication fills the context window and must be economized.
  - Compaction of context is lossy.
- Task decomposition represents the division of labor.
- The question of optimal scaffolding has natural analogues in:

  - Operations

  - Mechanism design

  - Industrial organization

Roadmap

Hayekian reasons for markets

- Distributed knowledge about:
    - Capabilities.
    - Preferences.
    - Constraints.

- Leads to decentralized mechanisms.

- Prices arise as an efficient way of communicating this decentralized information.

The American Economic Review

VOLUME XXXV SEPTEMBER, 1945 NUMBER FOUR

THE USE OF KNOWLEDGE IN SOCIETY

By F. A. HAYEK*

I

What is the problem we wish to solve when we try to construct a rational economic order?

On certain familiar assumptions the answer is simple enough. *If* we possess all the relevant information, *if* we can start out from a given system of preferences and *if* we command complete knowledge of available means, the problem which remains is purely one of logic. That is, the answer to the question of what is the best use of the available means is implicit in our assumptions. The conditions which the solution of this optimum problem must satisfy have been fully worked out and can be stated best in mathematical form: put at their briefest, they are that the marginal rates of substitution between any two commodities or factors must be the same in all their different uses.

This, however, is emphatically *not* the economic problem which society faces. And the economic calculus which we have developed to solve this logical problem, though an important step toward the solution of the economic problem of society, does not yet provide an answer to it. The reason for this is that the "data" from which the economic calculus starts are never for the whole society "given" to a single mind which could work out the implications, and can never be so given.

The peculiar character of the problem of a rational economic order is determined precisely by the fact that the knowledge of the circumstances of which we must make use never exists in concentrated or integrated form, but solely as the dispersed bits of incomplete and frequently contradictory knowledge which all the separate individuals possess. The economic problem of society is thus not merely a problem

* The author is Tooke professor of political economy and statistics at the University of London (London School of Economics and Political Science).

Market for contract labor

- Workers bid on jobs.

- In order to bid, workers must have information about costs and ability to complete the task.

- Mistakes result in cost overruns, loss to reputation, or removal from the platform.

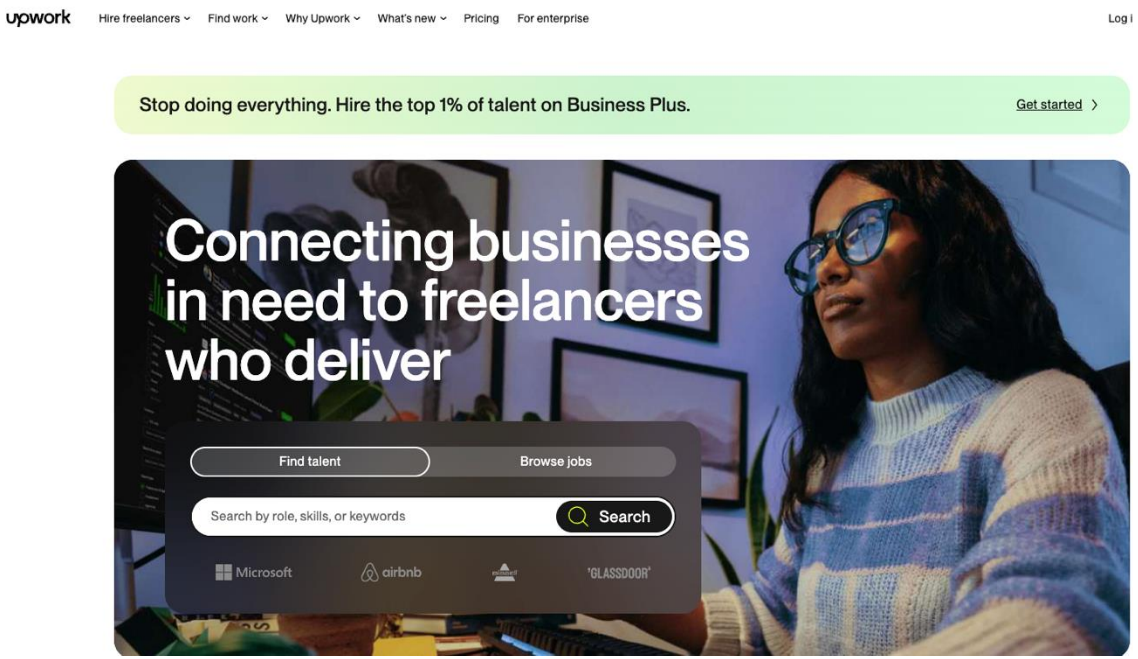

What is an agent?

- For loop?
- Combination of models + compute environment + budget + context + owner.
- Current conception of an agent is a specific instance of an LLM is too thin.

Market analogy for AI agents

- What are agents' costs? The number of tokens used in completing the task.

- What is success? Does the agent succeed in completing the task.

- Do agents have private information? Perhaps, if they have context or if they can "introspect."

- What is reputation? Doesn't exist for each unique instance.

- What are preferences? Unclear.

Production functions for AI agents

- An agent has many ways to do something.
- An agent may be able to choose among bundles of probability of success and tokens used.
- Becomes more complex when tools are available (including calls to different instances of a model).
- We don't have an understanding of which one is chosen.

Roadmap

1. How existing scaffolding for AI agents works.
2. Why a market-based scaffolding is appealing.
3. **Experiments with market-based scaffolding.**
4. Market design and concluding thoughts.

Using markets to coordinate agents?

- Need agents to know something about costs and probability to succeed.
- Need agents to pick appropriate production plans.
- Need agents to respond to price signals.

Testing the calibration of models

We evaluated 6 frontier models on 93 SWE-bench Lite tasks.

Tasks are GitHub issue and pull request pairs.

The patch is evaluated by running a set of tests.

SWE-bench

vals.ai

UPDATED: 2/24/2026

Solving production software engineering tasks

View all | Models only | Products only

All Models

Task type : Overall

	Systems (55)	Accuracy	Cost / Test	Latency	Settings
1	Claude Opus 4.6 (Thinking)	79.20 % ± 1.82	$1.69	414.36 s	
2	Gemini 3 Flash (12/25)	76.20 % ± 1.91	$0.41	426.90 s	
3	Claude Sonnet 4.6	76.20 % ± 1.91	$1.44	539.25 s	
4	GPT 5.2	75.40 % ± 1.93	$1.95	741.57 s	
5	GPT 5.3 Codex	75.20 % ± 1.93	$0.93	587.74 s	
6	Claude Opus 4.5 (Nonthinking)	74.60 % ± 1.95	$1.70	305.95 s	
7	Claude Opus 4.5 (Thinking)	74.20 % ± 1.96	$1.57	441.38 s	
8	Gemini 3 Pro (11/25)	71.60 % ± 2.02	$1.78	674.11 s	
9	GPT 5.1 Codex Max	71.00 % ± 2.03	$1.56	1739.97 s	
10	MiniMax-M2.5	70.40 % ± 2.04	$0.40	420.11 s	

Models are poorly calibrated

Model	Mean p	Pass rate	Skill	Est. toks	Realized toks
Claude Opus 4.5	76.5%	80.6%	+0.060	14,774	35,820
Gemini 3 Pro Preview	92.9%	80.6%	-0.111	1,801	55,969
GPT-5.2	63.2%	80.6%	-0.195	2,909	25,036
GPT-5.2-pro	66.3%	79.6%	-0.111	2,647	3,970
Claude Sonnet 4.5	75.8%	77.4%	+0.018	18,333	53,085
GPT-5-mini	61.4%	75.3%	-0.305	2,595	28,276

Single-model auction with reserve price

- Suppose there is an auction with a random reserve price that is not known to the agent.

- The agent bids to complete the task (or can choose not to bid) using an expected value calculation based on the estimated token cost and probability of success.

- If the agent wins the auction but fails to complete the task, it loses the auction and the dollar amount.

- If the agent wins the auction and completes the task, it gets the dollar amount minus the token cost.

Models are far from oracle when bidding

Model	*N*	Win rate	Exp. profit	Realized profit	Oracle profit
Claude Opus 4.5	93	28.2%	$0.077	$0.080	$0.260
Claude Sonnet 4.5	93	29.3%	$0.062	$0.064	$0.288
Gemini 3 Pro Preview	93	84.6%	$0.353	$0.303	$0.391
GPT-5-mini	93	21.6%	$0.054	$0.050	$0.375
GPT-5.2	93	5.4%	$0.005	$0.006	$0.385
GPT-5.2-pro	93	5.3%	$0.007	$0.007	$0.227

Market based scaffolding for AI agents

- There is a queue of tasks to be completed.
- Each workers submits a bid for a task.
- A bid is an ask, self assessment of probability of success, and expected completion time (currently not used).
- For each bid, the engine computes: $\text{Score}_i = p_i \times (\text{Utility} - \text{Ask}_i) - (1 - p_i) \times \text{Penalty}_i - E[\text{Cost}_i]$.
- All bids with Score >= 0 are considered, and the one with the highest score is awarded the task first unless otherwise occupied.
- If failure, next highest score is considered.

Some early results

Execution paradigm	Pass rate	Passes
Market (our scaffold)	58.0%	29/50
Solo GPT-5.2 (our scaffold)	48.0%	24/50
Solo GPT-5.2 (external scaffold)	74.0%	37/50
Oracle ceiling (external scaffold)	84.0%	42/50

So what?

- The design space is vast and interacts with model self-knowledge.
- What sort of bidding language is needed?
- Does it make sense to have a 'reputation' or alternatively an ML based router on top of this?
- How much better can this do vs the best designed scaffolding?

Roadmap

1. How existing scaffolding for AI agents works.
2. Why a market-based scaffolding is appealing.
3. Experiments with market-based scaffolding.
4. **Market design and concluding thoughts.**

Building good agents will become a key topic

- Learning the principal's preferences efficiently
- Learning what decisions they can make on their own and which they cannot
- Making them rational (no money pumps)
- Making them resistant to manipulation

The good place

- Improved preference elicitation and lower search costs improve taste matching.
- Better matching strengthens incentives for firms to invest in quality.
- Advertising may involve less puffery, especially with stronger reputation systems.
- In bargaining, AI lowers the opportunity cost of negotiating; negotiations can start earlier and persist longer, leading to better outcomes.
- Firms have less scope to extract rents from behavioral biases and bounded rationality.

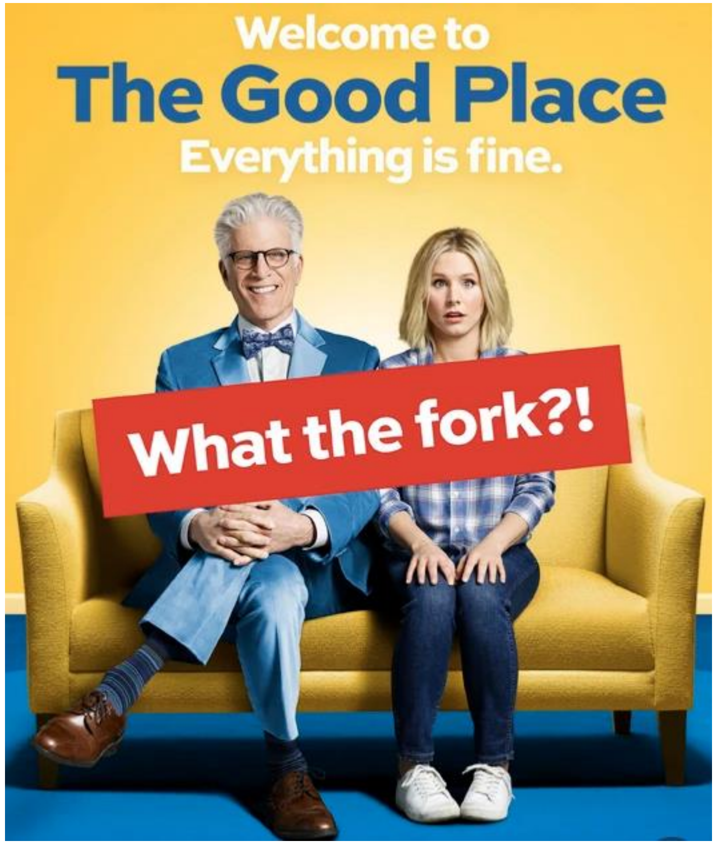

The bad place, ripoff hell

- Prices and dispersion may not necessarily go down.
- Firms may use obfuscation techniques to preserve or increase rents.
- Markets with externalities: e.g., congestion in AI-assisted job applications
- Zero-sum games may lead to race to the bottom, or arms race to exploit rents from having a “better” agent.

Getting to the good place with better mechanisms

Matching:

- Deferred acceptance at scale (e.g., dating apps, online labor markets)
- Costly signaling:
  - Journal submission, job applications, etc. need to be costly.
- Proof of humanity and proof of agent ownership.
  - Need to be able to prove that you are human and that you own the agent.
- Congestion taxes and rationing mechanisms.
  - Cloudflare's pay per crawl.
- Bring your own agent.
  - Acts on your behalf and helps you not to be exploited.
- Clear liability rules.

Conclusion

- Market mechanisms will be useful in coordinating AI agents.
- Agents require self-knowledge to use markets effectively.
- Simultaneously, we need to be thinking about how to develop market mechanisms to coordinate agents and prevent bad outcomes due to market failures.

Thank you

Andrey Fradkin, afradkin@gmail.com

Comments welcome