Ask a Buddy or a Bot?

Unbundling Answers from Relationships

Nicole Immorlica
Microsoft Research
& Yale University

Brendan Lucier
Microsoft Research

Joseph E. Stiglitz
Columbia University
& NBER

Màxim Ventura-Bolet
Columbia University

Network Science & Economics, University of Luxembourg

October 25, 2026

This paper

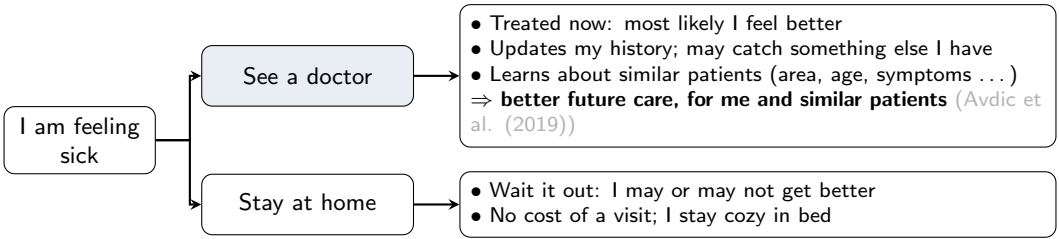
Can AI chatbots reduce the number of questions we ask?
Can they reduce the value of the answers we receive?

We build a simple model of rational agents in which:

1. **Short run:** AI weakly improves both the number of questions answered and the value of the answers.
2. **Medium run:** displacing human consultations triggers a negative feedback loop through erosion of social capital.
3. **Long run:** welfare can end up *lower* than before AI — possibly both fewer questions asked and less valuable answers.

Motivation

I am feeling sick. How can I get better?



Routing can be inefficient: the private value of seeing the doctor is below its social value (I ignore benefit to similar patients).

AI chatbots appear

New option: describe my symptoms to an **AI chatbot**. I ask it if:

$$\text{Chatbot's advice} \geq \max\Big\{\underbrace{\text{Doctor's treatment} + \text{my better future care}}_{\text{again ignoring other similar patients}},\ \text{staying at home}\Big\}$$

- Today the choice set grows so we are weakly better off
- Tomorrow's options depend on today's actions... "see a doctor" gets worse.

AI chatbots appear

New option: describe my symptoms to an **AI chatbot**. I ask it if:

$$\text{Chatbot's advice} \geq \max\Big\{\underbrace{\text{Doctor's treatment} + \text{my better future care}}_{\text{again ignoring other similar patients}},\ \text{staying at home}\Big\}$$

- Today the choice set grows so we are weakly better off
- Tomorrow's options depend on today's actions... "see a doctor" gets worse.

Two features of the example (agnostic about *why* in both cases)

- Questions are *heterogeneous*.
- Questions are *related*.

Remarks

Theoretical thought experiments are helpful here:

- Past shocks may be different.
- Empirics can't capture equilibrium effects.

Related work:

- AI reorganizes knowledge hierarchies (Ide & Talamàs 2025); changes information-production incentives (Stiglitz & Ventura-Bolet 2025); disrupts tacit-knowledge transmission (Ide 2026); lowers learning incentives (Acemoglu, Kong & Ozdaglar 2026); erodes workers' skills (Caosun & Aral 2026).

Today:
AI changes how questions are *routed*, which endogenously changes the future value of human consultation.

Plan

1. **Pre-AI answering system**
   Question routing, social capital, and steady states
2. **Impact of AI**
   Short run $\rightarrow$ Medium run $\rightarrow$ Long run
3. **Social Planner**
   The dynamic externality and the planner's problem
4. **Discussion**

Pre-AI answering systems

- Time is discrete $t = 0, 1, 2, ...$
- Mass M of atomistic players
- Each period, each player receives one question $m \in [0, M]$
- She can consult another human or not:

$$\underbrace{u^H(m, S_t)}_{\text{consult a human}} \quad \text{vs.} \quad \underbrace{u^N(m)}_{\text{do not consult}}$$

Gain from consulting a human:

$$\boxed{G(m, S_t) \equiv u^H(m, S_t) - u^N(m)}$$

Assumptions

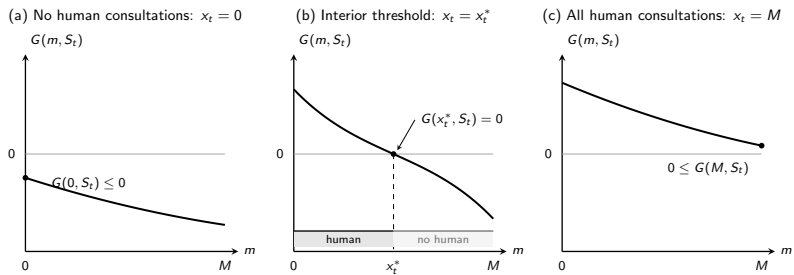
$$G_m < 0, \qquad G_S > 0, \qquad G_{mS} \geq 0$$

- $G_m < 0$: questions ordered by gain from a human
- $G_S > 0$: social capital improves human answers
- $G_{mS} \geq 0$: changes in S preserve the ordering

Pre-AI: how to route questions?

Players are atomistic: own choice does not move aggregate $S \Rightarrow$ each optimizes period by period. Measure of questions asked to humans:

$$x_t = x(S_t) = \int_0^M \mathbf{1}\{G(m, S_t) \geq 0\}\, dm \qquad (1)$$



Pre-AI: equilibrium path

Social capital evolves as:

$$\boxed{S_{t+1} = (1-\delta)S_t + I(x_t)} \qquad (2)$$

$$\delta \in (0,1), \qquad I_x > 0$$

more human consultations today

$$\Downarrow$$

$S_{t+1} \uparrow$

$$\Downarrow$$

better human answers tomorrow

Given an initial condition S_0, equations (1) and (2) jointly determine the equilibrium path $\{x_t, S_t\}_{t \geq 0}$.

Individual choices are static, but aggregate routing generates dynamics through social capital.

Pre-AI: welfare

Flow welfare:

$$W(S_t) \equiv \int_{G(m,S_t)\geq 0} u^H(m, S_t)\, dm + \int_{G(m,S_t)<0} u^N(m)\, dm$$

Lifetime welfare:

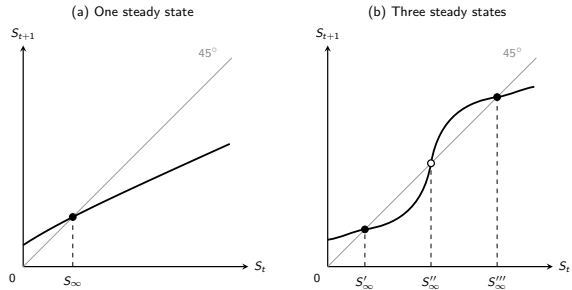
$$V(\{S_t\}_{t=0}^{\infty}) \equiv \sum_{t=0}^{\infty} \beta^t W(S_t)$$

Observation

W is increasing in $S \Rightarrow$ a path with uniformly higher social capital gives higher lifetime welfare.

Pre-AI: steady states

Steady state: $x_t = x(S_\infty)$ and $S_\infty = (1-\delta)S_\infty + I(x(S_\infty))$.



If $I(0) > 0$ and $G\left(M, \frac{I(M)}{\delta}\right) < 0$: odd number of steady states, all interior; stable and unstable steady states alternate.

AI arrives at a pre-AI steady state $S^{pre\text{-}AI}$: some questions asked to humans, some left unanswered.

AI chatbots arrive: new equilibrium

A new option: ask the AI chatbot, with payoff $u^{AI}(m)$.

Gain from consulting a human now compares against the best alternative:

$$G^{AI}(m, S_t) \equiv u^H(m, S_t) - \max\{u^N(m), u^{AI}(m)\}$$

$$\boxed{x^{AI}(S_t) = \int_0^M \mathbf{1}\{G^{AI}(m, S_t) \geq 0\}\, dm} \qquad (1^*)$$

Everything else is unchanged: social capital still evolves by (2). Given $S_0 = S^{pre\text{-}AI}$, equations (1^*) and (2) determine the equilibrium path $\{x_t^{AI}, S_t\}_{t \geq 0}$.

AI chatbots arrive: short-run impact

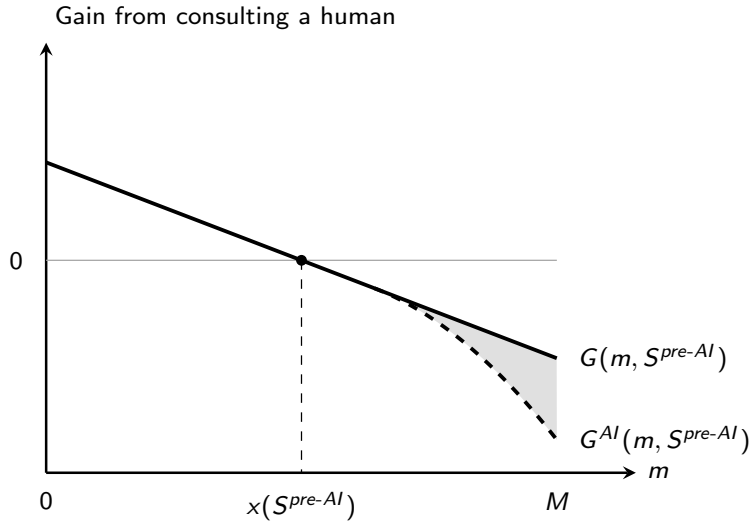
Short-run welfare gain

$$W^{AI}(S^{\text{pre-AI}}) - W(S^{\text{pre-AI}}) = \underbrace{\int_{G(m,S^{\text{pre-AI}})\geq 0} \max\{0, u^{AI}(m) - u^{H}(m, S^{\text{pre-AI}})\}\, dm}_{\text{better answers to previously answered questions}}$$

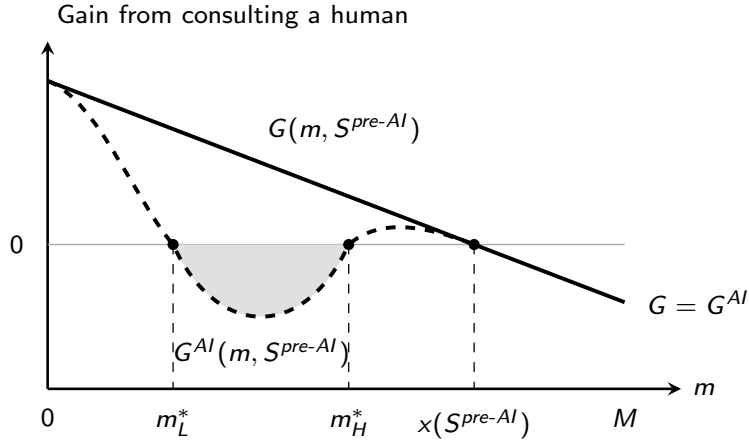
$$+ \underbrace{\int_{G(m,S^{\text{pre-AI}})<0} \max\{0, u^{AI}(m) - u^{N}(m)\}\, dm}_{\text{answers to previously unanswered questions}} \geq 0$$

Both give immediate gains; only substitution of humans starts the social-capital dynamics.

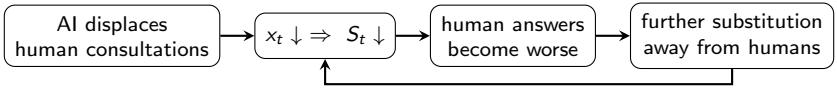
Short run: AI answers previously unanswered questions

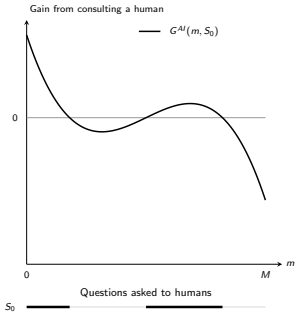


Short run: AI substitutes for human answers



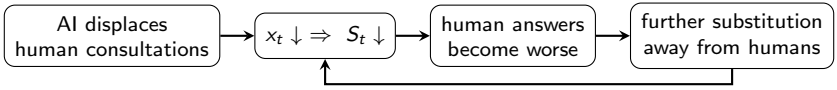
Medium run: the human answering system adjusts

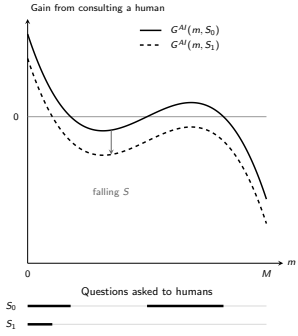




- $S' \leq S \Rightarrow x^{AI}(S') \leq x^{AI}(S)$, if AI initially displaces a positive measure of human consultations, $S_{t+1} < S_t$ and $x_{t+1} \leq x_t$ for all t, $W^{AI}(S_{t+1}) \leq W^{AI}(S_t)$ for all t.

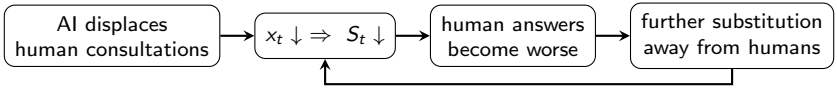
Medium run: the human answering system adjusts

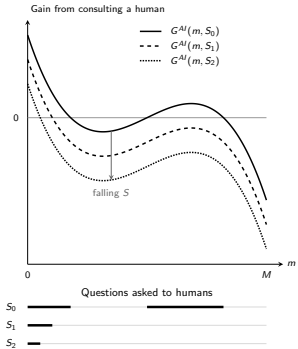




- $S' \leq S \Rightarrow x^{AI}(S') \leq x^{AI}(S)$, if AI initially displaces a positive measure of human consultations, $S_{t+1} < S_t$ and $x_{t+1} \leq x_t$ for all t, $W^{AI}(S_{t+1}) \leq W^{AI}(S_t)$ for all t.

Medium run: the human answering system adjusts





- $S' \leq S \Rightarrow x^{AI}(S') \leq x^{AI}(S)$, if AI initially displaces a positive measure of human consultations, $S_{t+1} < S_t$ and $x_{t+1} \leq x_t$ for all t, $W^{AI}(S_{t+1}) \leq W^{AI}(S_t)$ for all t.

Long run: are we still getting better answers?

Convergence to the largest AI steady state below $S^{\text{pre-AI}}$: $\delta S_\infty^{AI} = I\left(x^{AI}(S_\infty^{AI})\right)$, with

$$W^{AI}(S_\infty^{AI}) = \int_0^M \max\left\{u^H(m, S_\infty^{AI}),\ u^{AI}(m),\ u^N(m)\right\} dm$$

$$W^{AI}(S_\infty^{AI}) - W(S^{\text{pre-AI}}) = \text{impact gain} - \text{erosion on diverted questions} - \text{erosion on questions still asked to humans} \gtrless 0$$

Bottomline: it depends.

Long run: are we still getting better answers?

Convergence to the largest AI steady state below $S^{\text{pre-AI}}$: $\delta S_\infty^{AI} = I\left(x^{AI}(S_\infty^{AI})\right)$, with

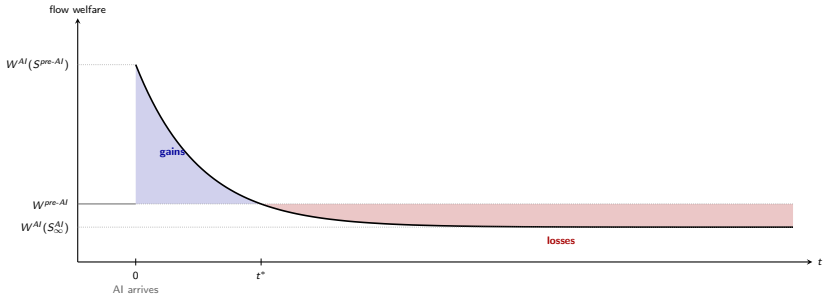
$$W^{AI}(S_\infty^{AI}) = \int_0^M \max\left\{u^H(m, S_\infty^{AI}),\ u^{AI}(m),\ u^N(m)\right\} dm$$

$$W^{AI}(S_\infty^{AI}) - W(S^{\text{pre-AI}}) = \text{impact gain} - \text{erosion on diverted questions} - \text{erosion on questions still asked to humans} \gtrless 0$$

Bottomline: it depends.

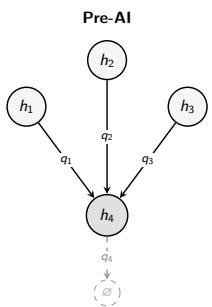
- One could think that this depends on an underlying network structure...

Example:

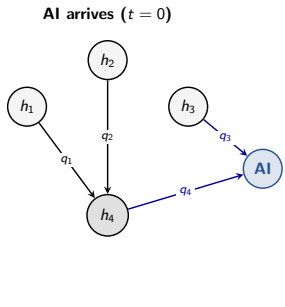


- Gains until t^*, then losses in every period thereafter
- $V^{AI} \gtrless V^{pre\text{-}AI}$ depends on how β weighs the two areas: impatient societies gain, patient ones lose

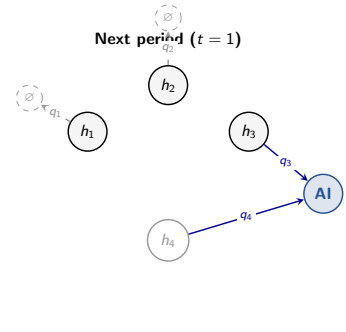
Long run: will we ask more questions?



h_4 (doctor) answers 3 questions, S high own question q_4 unanswered
3 answered



AI takes q_3 from the doctor and answers q_4, previously unanswered
4 answered ↑



S fell: h_4's answers no longer worth asking
AI can't do q_1, q_2
2 answered ↓ (below pre-AI)

AI can initially make society ask more questions and eventually leave it asking fewer.

The Social Planner

Users pick the best current answer. The planner also values the future social capital created by human consultations.

Let $\mathcal{H}_t \subseteq [0, M]$ be the set of questions assigned to humans:

$$\max_{\{\mathcal{H}_t\}_{t=0}^{\infty}} \quad \sum_{t=0}^{\infty} \beta^t \left[\int_{\mathcal{H}_t} u^H(m, S_t)\, dm + \int_{[0,M]\setminus\mathcal{H}_t} \max\{u^N(m), u^{AI}(m)\}\, dm \right]$$

$$\text{s.t.} \quad S_{t+1} = (1-\delta)S_t + I(\mu(\mathcal{H}_t)), \qquad S_0 = S^{\text{pre-AI}}$$

Given $\mu(\mathcal{H}_t)$, the planner assigns to humans the questions with the highest $G^{AI}(m, S_t)$.

What we know

With $\mathcal{H}^P(S)$ the planner's optimal consultation set, starting from the same $S^{\text{pre-AI}}$:

$$\mu(\mathcal{H}^P(S)) \geq x^{AI}(S), \qquad S_t^P \geq S_t$$

The planner preserves weakly more human interaction and social capital, potentially sacrificing current payoff to improve future welfare.

Discussion: the network behind social capital

After two days hearing about key players, cascades, endogenous network formation, dynamic networks...

1. Social capital: a stock of past human consultations (network activity). Can an explicit network layer provide a foundation for our baseline model?
2. Different questions may contribute differently to future answering capacity. At the same total volume of human consultations, does it matter who asks whom and which questions they ask?
3. Which questions, when diverted to AI, make other human consultations no longer worthwhile? Can the loss of particular questions, experts, or links trigger a cascade and collapse of the human answering network?

Thank you!