What You Don’t Know May Be Good For You

Johannes Hörner and Larry Samuelson

Yale University

University of Nevada, Reno
November 7, 2025

I. Introduction

I.1 Motivation

- In the late 1980s, the state of New York introduced its Cardiac Surgery Reporting System.
  - A surgeon's assessment: "The so-called best surgeons are only doing the most straightforward cases."
  - The economists' assessment: "Mandatory reporting mechanisms inevitably give providers the incentive to decline to treat more difficult and complicated patients." (Dranove, Kessler, McClellan and Satterthwaite, 2003)
  - The general phenomenon—Goodhart's law: "When a measure becomes a target, it ceases to be a good measure."
- Our question: How are reputations built when the builder can manipulate the relevant information?

I. Introduction

I.1 Motivation

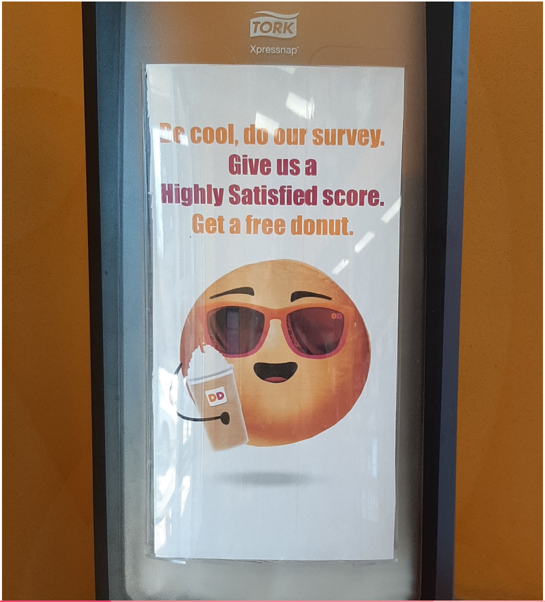

I.2 This paper

- We examine an interaction between a long-run expert and a series of short-run clients.
- In each “period”, the expert selects a type of client, serves that client, and generates an outcome.
- In the benchmark, the type of client and the outcome are unobserved.
- A report card makes the outcome observable, but not the type of client.

I. Introduction

I.3 Preview

- Benchmark questions: What is the equilibrium outcome and what is the efficient allocation in
  - the stage game,
  - the repeated game without report cards, and
  - the repeated game of *perfect* monitoring?
- The equilibrium outcome in the repeated game (of imperfect monitoring) with report cards has
  - some interesting structure,
  - and some inefficiencies that can overwhelm the information benefits.
- And, time permitting, various extensions.

- Probabilities of success are

		Expert B	Expert G
Client	L	1	1
	H	$1-z$	1

.

- Given probability q of the expert being good, expected payoffs are:

$$
\begin{aligned}
\text{Low client :} \quad & 1-\phi_L &=& \; 0 \\
\text{High client :} \quad & q+(1-q)(1-z)-\phi_H &=& \; zq-z-\phi_H \\
& &=& \; zq-\theta .
\end{aligned}
$$

II. The Model

II.1 The stage game

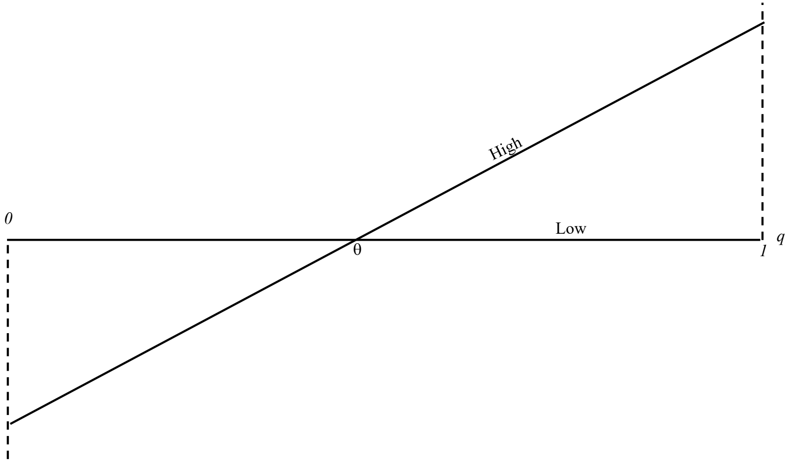

II. The Model

II.2 Stage-game equilibrium and efficient allocation

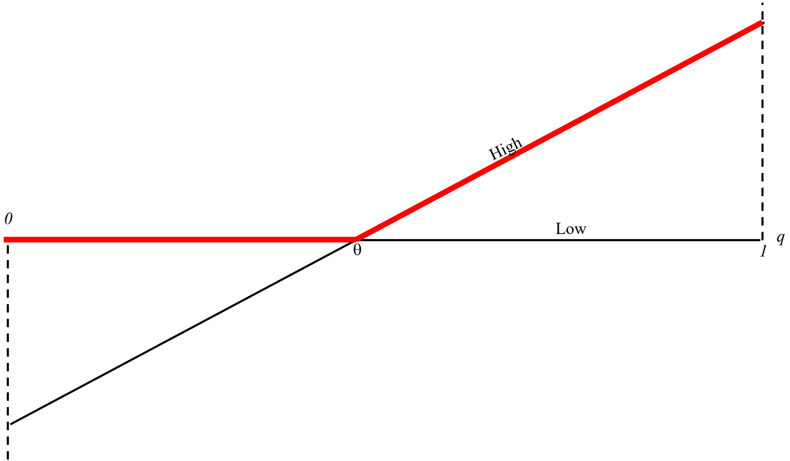

II. The Model

II.3 Repeated game (without report cards) equilibrium

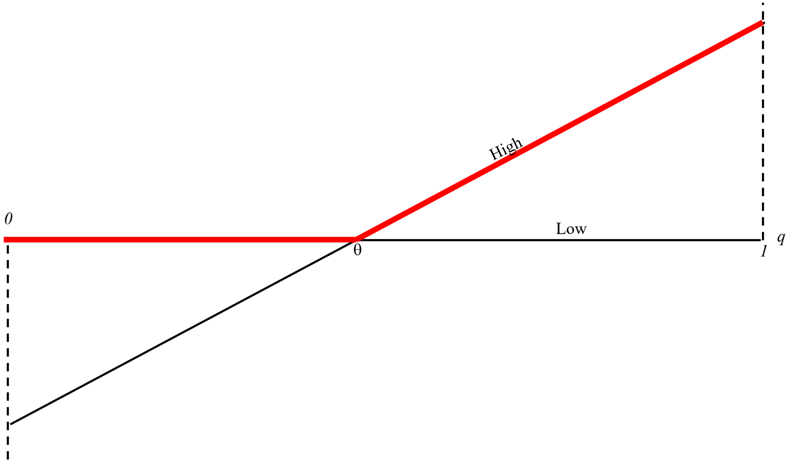

II. The Model

II.4 Repeated game of perfect monitoring, efficient allocation

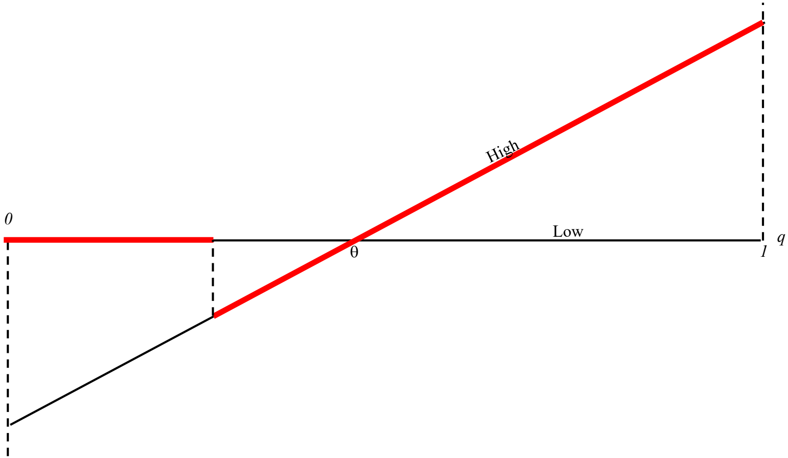

II. The Model

II.5 Repeated game of perfect monitoring, equilibrium

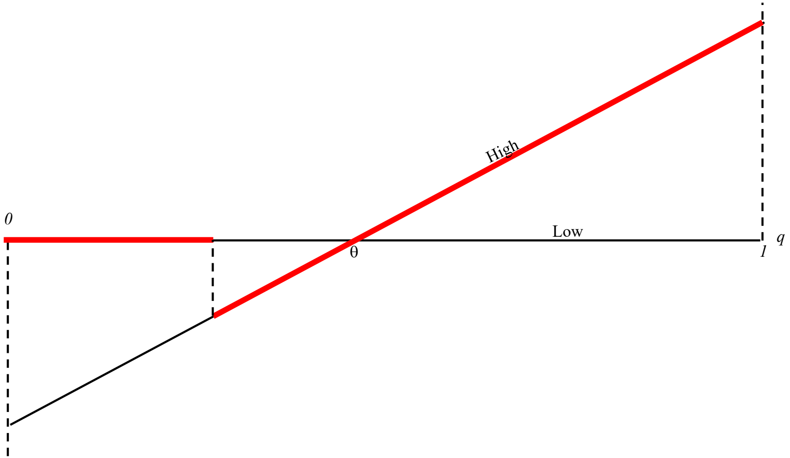

III. The Repeated Game with Report Cards

III.1 Structure

- Time is continuous, horizon is infinite.
- Continuum of experts, long-lived.
- Each expert has a type, G, B, prior $q^0 = \mathbf{P}[\omega = G]$.
- Continuum of clients (the "market"), short-lived, two types, H, L.

- The expert chooses the fraction $k_t \in [0,1]$ of H-clients.
- The path (k_t) defines a Poisson process $\{N_t : t \geq 0\}$ with intensity

$$\lambda(t) := \lambda \mathbf{1}_{\{\omega=B\}} k_t.$$

- The price is $P_t = q_t - \theta$ for high clients, zero for low clients.
- Discount rate r.

- Given the path $(k_s)_{s\leq t}$, the expert's belief about his type at time t is given by
$$\tilde{q}(t) = \mathbf{P}\left(\omega = G \mid \sigma\left(N_s, k_s \colon s \leq t\right)\right).$$
- The market belief is q.
- $q = 1 \implies$ high clients.
- $q = 0 \implies$ low clients.
- Focus on experts with no failures. Let $\tau \in \mathbf{R}_+ \cup \{+\infty\}$ be the time of the first arrival.

III. The Repeated Game with Report Cards

III.2 Strategies

- A deterministic policy is a measurable function $\pi : \mathbf{R}_+ \to [0,1]$, which specifies k_t conditional on $\{t < \tau\}$.

- A deterministic policy is
  - extreme-valued if $k_t \in \{0,1\}$ for all t;
  - a stopping policy if there is s such that

$$\pi_t(s) = \begin{cases} 0 \text{ for } s < t, \\ 1 \text{ for } s \geq t; \end{cases}$$

  - a cutoff policy if there is $\bar{q}$ such that

$$\pi_t(s) = \begin{cases} 0 \text{ if } q_s < \bar{q}, \\ 1 \text{ if } q_s \geq \bar{q}; \end{cases}$$

  - a Markov policy if $q_t = q_s \implies \pi(t) = \pi(s)$.
- A mixed policy is a mixture over deterministic policies.

III. The Repeated Game with Report Cards

III.3 Beliefs

- Let

$$\mathrm{d}H_t(q) = \mathbf{P}^{\phi}\left(\tilde{q}_t = q \mid \sigma\left(N_s : s \leq t\right)\right)$$

  be the belief of a (selected) H client *about the expert's belief*.
- Let q_t denote its mean.
- If the expert's policy is not deterministic, $q_t \neq \tilde{q}_t$, the expert's own belief.

- An *equilibrium* is a pair (H, ϕ) (and hence q) such that:

  1. Given (q_t), ϕ maximizes

  $$\mathbf{E}^{\phi}\left[\int_{t\geq 0} re^{-rt} k_t(q_t - \theta)\mathrm{d}t \mid \tilde{q}_0 = q^0\right];$$

  2. Given ϕ, it holds that, for all $t \geq 0$, all $q \in [0,1]$,

  $$\mathrm{d}H_t(q) = \mathbf{P}^{\phi}\left(\tilde{q}_t = q \mid \sigma\left(N_s : s \leq t\right)\right).$$

- Also, assume $q_0 = q^0$.

- The efficient policy is a cutoff policy,

$$k_t = 1 \text{ iff } q_t \geq \bar{q},$$

for some $\bar{q} \in (0, \theta)$.

- The cutoff is determined by

$$q \lesseqgtr \bar{q} \Leftrightarrow 0 \gtreqless \int_{t \geq 0} e^{-rt} \left(q + (1-q)e^{-\lambda t}\right) \left(\frac{q}{q + (1-q)e^{-\lambda t}} - \theta\right) \mathrm{d}t.$$

- This is also the equilibrium outcome when monitoring is perfect.

III. The Repeated Game with Report Cards

III.5 Efficient allocation

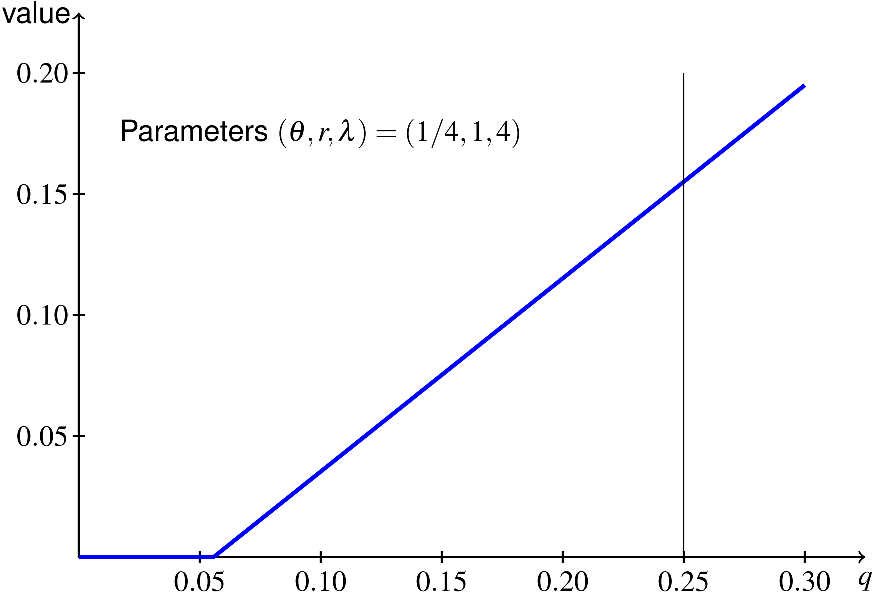

III. The Repeated Game with Report Cards

III.6 Equilibrium is inefficient

- The efficient policy is not an equilibrium policy - no cutoff policy is.
  - Suppose $\bar{q} < \theta$. Consider $q^0 \in (\bar{q}, \theta)$:
    - $q^0 - \theta < 0$: taking L-clients yields a higher reward that H clients;
    - Taking L-clients guarantees that no failure occurs.
    - Hence, for $q^0 < \theta$, the expert must take L-clients in equilibrium.
  - Suppose $\bar{q} > \theta$. Then an expert with $q^0 \in (\theta, \bar{q})$ will take H rather than L client.
  - So $\bar{q}$ must equal θ. Then the argument of the first case ensures that experts with q^0 just above θ will take L rather than H clients.

III. The Repeated Game with Report Cards

III.6 Equilibrium is inefficient

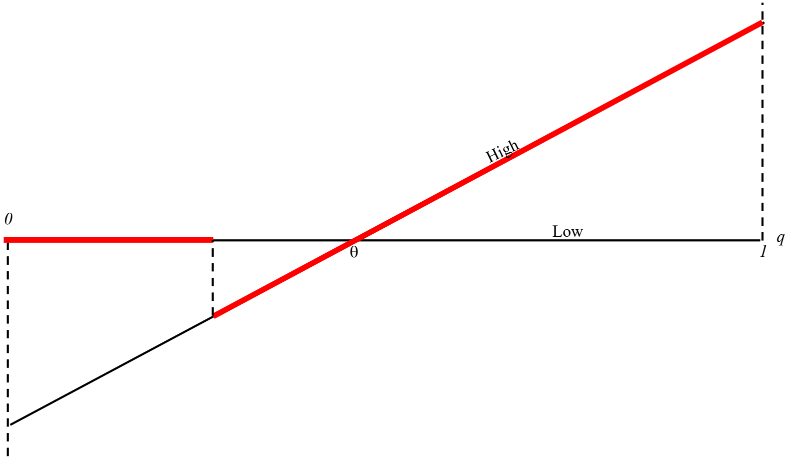

III. The Repeated Game with Report Cards

III.6 Equilibrium is inefficient

III. The Repeated Game with Report Cards

III.7 The temptation to wait

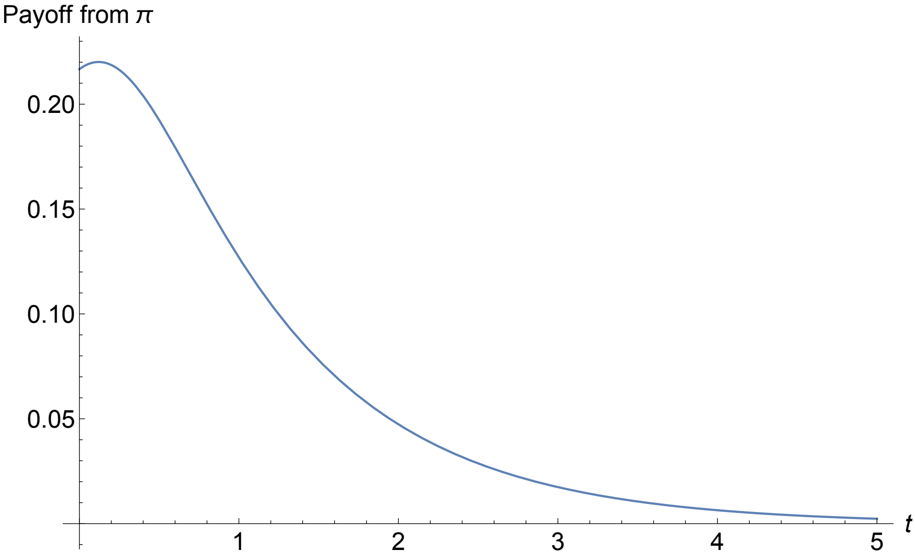

III. The Repeated Game with Report Cards

III.8 Experts must mix, but it's not clear they can

- To see that the equilibrium must be mixed, note that:
  - Consider $q^0 > \theta$ ("not too close to 1").
  - If the expert is expected to take only L clients, then the expert will prefer to take H clients, which are then more profitable.
  - If the expert is expected to take H-clients *only*, the price increases "fast." Then it is profitable to delay H clients.
- The difficulty: clients learn too fast and prices increase too fast.
- The obvious way to use "mixtures" to slow learning is for experts to choose $k \in (0, 1)$, but this will not work.
- Suppose the expert is expected to take *a fraction* of H-clients.
- By deviating to H-clients only (say), the expert's private belief is then higher than the market's, conditional on $\{\tau < t\}$.
- If she is indifferent on path, this "over"-optimism means that she strictly prefers taking H-clients, giving a profitable deviation.

Proposition 1

An equilibrium exists. There exists $\bar{q} \in (\theta, 1)$ such that:

- *if $q^0 \leq \theta$, experts take low clients always.*
- *If $q^0 \geq \bar{q}$, experts take high clients always.*
- *If $q^0 \in (\theta, \bar{q})$, there exists $T > 0$ such that experts pick a random time $s \leq T$ according to some atomless distribution with full support on $[0, T]$, and play according to π_s, the stopping policy that switches from low clients to high clients at time s.*

This equilibrium is the only one in which the market assigns positive probability that an expert without failures takes on some high clients.

III. The Repeated Game with Report Cards

III.9 Equilibrium exists, is unique, and is a mixture over stopping times

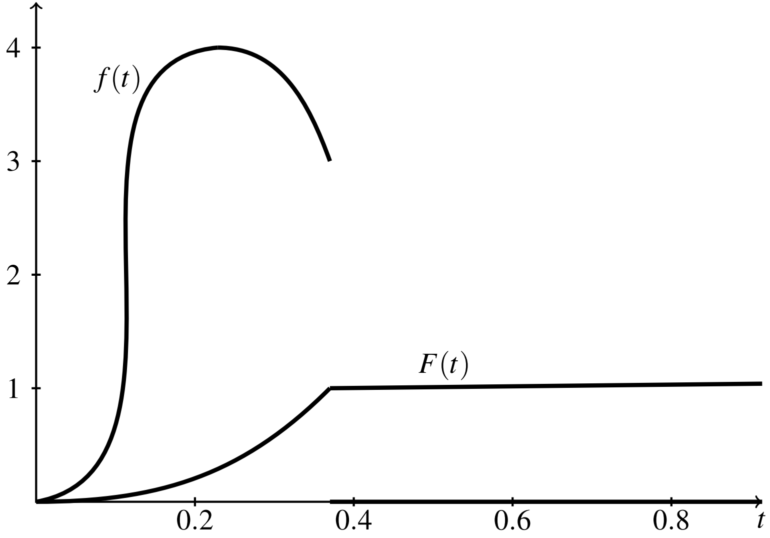

III. The Repeated Game with Report Cards

III.9 Equilibrium exists, is unique, and is a mixture over stopping times

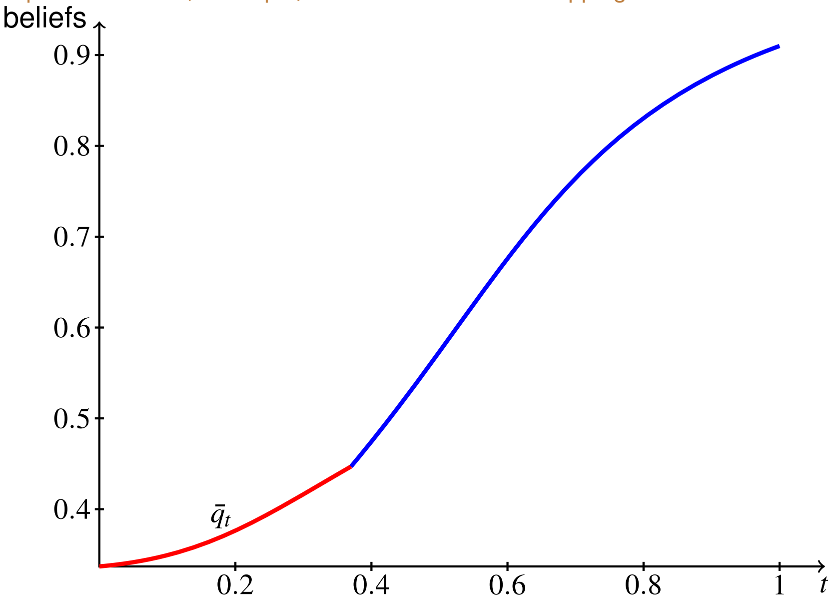

III. The Repeated Game with Report Cards

III.9 Equilibrium exists, is unique, and is a mixture over stopping times

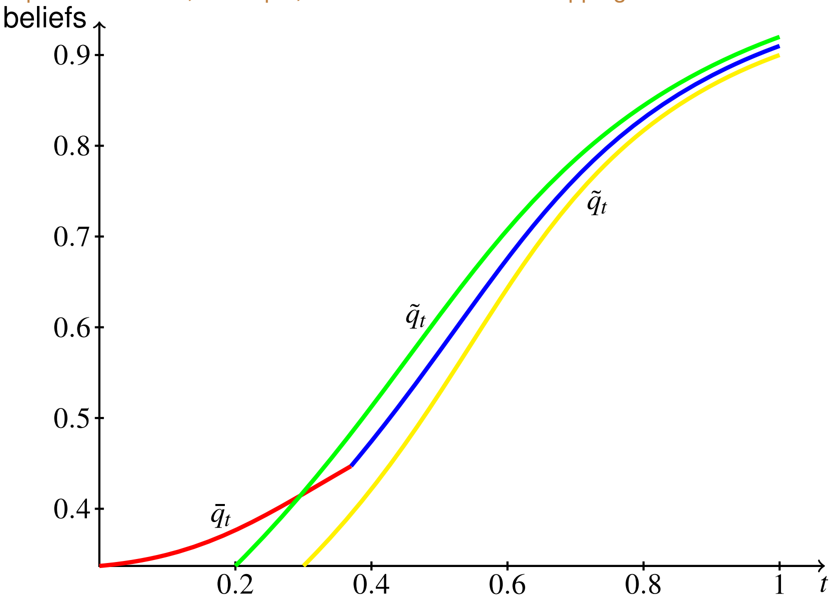

- The equilibrium with report cards
  - allows clients to learn about experts, and
  - ensures that any expert who has failed only low clients,
  - but experts sometimes inefficiently choose low clients even when $q_t > \theta$.
- The no-information solution
  - gives no learning, and
  - experts who have failed take high clients,
  - but experts choose high clients whenever $q_t > \theta$.
- Depending on parameter values, welfare may be higher in either of these equilibria.

IV. Extensions

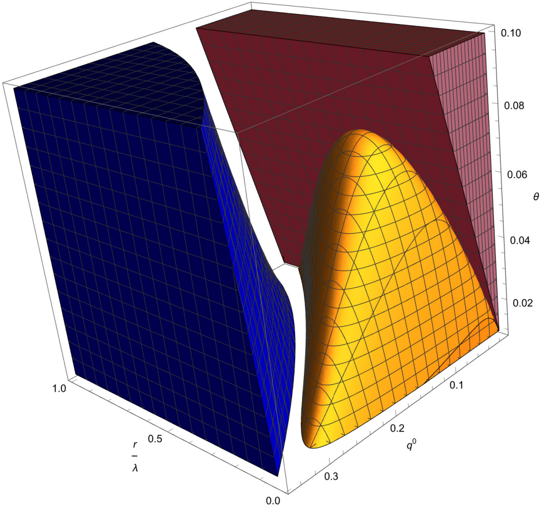

IV. Extensions

IV.1 The Tea Break Equilibrium

- An alternative, non-Markov equilibrium slows the growth of prices by scheduling “tea breaks,” in which only low clients are accepted, with experts otherwise accepting H clients.
- Accepting an H client during a tea break is deterred by adverse beliefs.
- Letting tea breaks be arbitrarily frequent and arbitrarily short, one can think of a density describing the fraction of time the expert takes high clients, as a function of time.
- This equilibrium produces lower welfare than the equilibrium with delay.

IV. Extensions

IV.1 The Tea Break Equilibrium

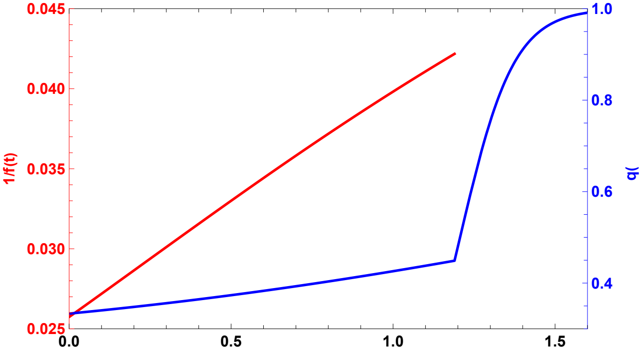

- Proportion of high clients taken (red) and price (blue).

IV. Extensions

IV.2 Randomly revealing lemons

- Suppose an expert experiencing a failure is revealed to have done so randomly, with hazard rate γ.
- $\gamma = 0$ is no information; $\gamma = \infty$ is report cards.
- We can set γ so that experts are willing to immediately and always take high clients.
- It is immediate that this scheme gives higher welfare than no information (and hence, sometimes, than report cards).

IV. Extensions

IV.3 Randomly hiding lemons

- Suppose failures are hidden with probability α_t. (Before, $\alpha_t = 1$.)
- Once hidden, a failure is randomly revealed with probability γ_t.
- This generalizes, and hence can improve upon the previous scheme and the no-information scheme.
- This also can improve upon the delay equilibrium.
  - Design α_t and γ_t to replicate the delay price path.
  - Confirm that experts will always take high clients.

IV. Extensions

IV.3 Randomly hiding lemons

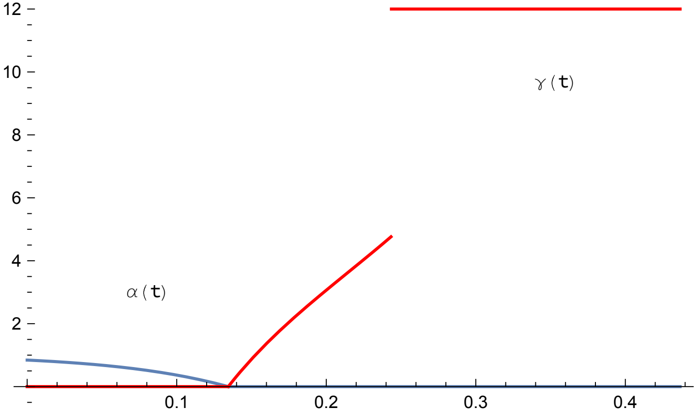

IV. Extensions

IV.4 Price Controls

- If prices were fixed and constant, then there would exist an efficient equilibrium.
- “Economists, incidentally, would have a ready solution to this entire problem: the inexperienced surgeon should reveal his inexperience but charge substantially less for the procedure, just the way numerous but inexperienced professionals from other professions do.”
- If types and outcome were observable, the complete competitive markets would be efficient.

IV. Extensions

IV.5 Good news

- Suppose all matches yield ordinary outcomes except good experts and high clients, who stochastically produce breakthroughs.
- The efficient outcome is a cutoff strategy, with beliefs drifting down toward the cutoff until (perhaps) a breakthrough pushes the posterior to one.
- The equilibrium with report cards will duplicate the efficient outcome.

IV.6 Unknown Client Type

- Suppose clients do not know their types, with a common prior of being high.
- Experts observe types,
- A client updates their belief, and hence willingness to pay, upon being accepted.
- Any client accepted by an expert of a given vintage, who has experienced no failure, must command the same price.
- In equilibrium, only low clients are accepted.

IV. Extensions

IV.7 Learning by doing

- The expert's type, known to the expert, is initially G with probability q^0.
- The expert's type changes from B to G with probability γk_t; type G is absorbing.
- The market observes outcomes, with failures arriving at rate λ in state B (only).
- The expert takes high clients in state G.
- Conditional on B, the expert's private history is irrelevant to his incentives and the expert's private belief remains degenerate.
- The B-expert initially diversifies his clientele, accepting an increasing fraction k_t of high clients, until some time at which he switches to high clients

IV. Extensions

IV.7 Learning by doing

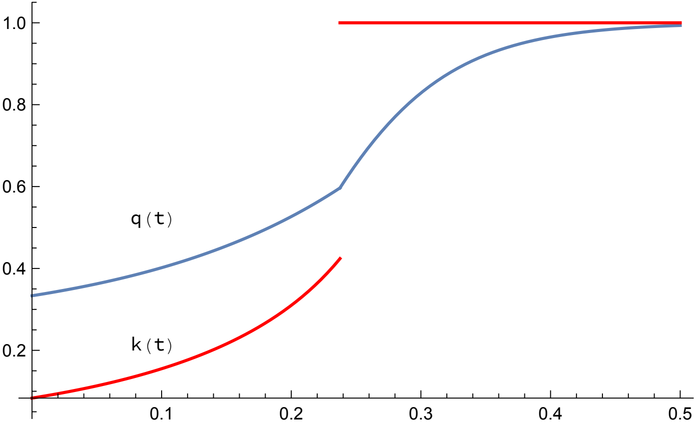

- q_t shows the belief about the expert's type. k_t shows the fraction of high clients taken by the expert at time t, which jumps to 1 after an initial period of a mixed clientele.

V. Discussion

THANK YOU