Lecture 10: Reputation in Markets

Alexander Wolitzky

MIT

14.126, Spring 2024

Reputation in Markets

Last class covered the canonical LR-SR reputation model and some extensions.

- Focus primarily on general equilibrium selection.

Today, embed reputation formation in more "market-like" models.

- Focus primarily on reputation as an ingredient in positive economic models.
- Often restrict attention to Markov perfect eqm. Problematic if goal is well-founded eqm selection, but fine for building a descriptive model.

Plan

1. Good reputation: reputation $\Longrightarrow$ incentives for **high** effort $\Longrightarrow$ better outcomes. (Holmström 82, Mailath-Samuelson 01, Board-Meyer-ter-Vehn 13)
2. Bad reputation: reputation $\Longrightarrow$ incentives for **wrong kind** of effort ("pandering") $\Longrightarrow$ worse outcomes. (Morris 01, Ely-Valimäki 03, Ely-Fudenberg-Levine 08)

Career Concerns (Holmström 82)

Holmström's "career concerns" model is a classic contract theory model that's also a reputation model.

An agent with ability $\theta \sim N(\mu_0, \tau_\theta)$ faces a labor market (population of potential employers). τ=precision=$1/\sigma^2$. No one observes θ.

- Different from most modern reputation models, where θ is the reputation-builder's private information. For tractability.

Each period t, agent chooses effort $a_t \geq 0$ at cost $c(a_t)$ ($c', c'' > 0$, $c'(0) = 0$).

Then everyone (agent and market) observe period t output

$$y_t = \theta + a_t + \varepsilon_t, \text{ where } \varepsilon \sim N(0, \tau_\varepsilon) \text{ iid.}$$

The agent is paid her expected output $\mathbb{E}[y_t | y^t]$ each period, so her payoff is

$$(1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} \left(\mathbb{E}[y_t] - c(a_t)\right).$$

Equilibrium

Let's find the (pure) eqm in Holmström's model.

- There can also be mixed equilibria, where since realized a_t is unobserved, the agent comes to hold different beliefs than the market. Holmström didn't consider this possibility. It still isn't well-understood.

When agent uses a pure strategy, given the history of output $y^t = (y_1, \ldots, y_t)$, the market can back out the history of on-path effort $a^{*t} = (a_1^*, \ldots, a_t^*)$, and hence also the history of output-net-of-effort $z^t = (z_1, \ldots, z_t)$, where

$$\begin{aligned} z_t &= y_t - a_t^* \\ &= \theta + \varepsilon_t + a_t - a_t^* =_{\text{in eqm}} \theta + \varepsilon_t. \end{aligned}$$

- Effort affects output, which affects market belief.
- Agent has positive return to effort, but market isn't fooled in eqm.

Equilibrium (cntd.)

In eqm, the z_t's are iid normal signals of a normal parameter.

A standard updating formula implies that the market's posterior belief conditional on z^t is that $\theta \sim N\left(\mu_{t+1}, \tau_{t+1}\right)$ with

$$\begin{aligned}
\mu_{t+1} &= \frac{\tau_t \mu_t + \tau_\varepsilon z_t}{\tau_t + \tau_\varepsilon} = \frac{\tau_\theta \mu_0 + \tau_\varepsilon \sum_{s=1}^t z_s}{\tau_\theta + t\tau_\varepsilon}, \\
\tau_{t+1} &= \tau_t + \tau_\varepsilon = \tau_\theta + t\tau_\varepsilon.
\end{aligned}$$

Note that $\mu_t \to \theta$ almost surely, and $\tau_t \to \infty$.

- Market eventually learns θ.

Note: Normality key for tractability.

Equilibrium (cntd.)

Let $a_t^*\left(y^{t-1}\right)$ denote agent's effort at history y^{t-1}. The agent's period t wage at history y^{t-1} is

$$\begin{aligned}\mathbb{E}\left[y_t|y^{t-1}\right] &= \mu_t\left(y^{t-1}\right)+a_t^*\left(y^{t-1}\right)\\ &= \frac{\tau_\theta}{\tau_t}\mu_0+\frac{\tau_\varepsilon}{\tau_t}\sum_{s=1}^{t-1}z_s+a_t^*\left(y^{t-1}\right)\\ &= \frac{\tau_\theta}{\tau_t}\mu_0+\frac{\tau_\varepsilon}{\tau_t}\sum_{s=1}^{t-1}\left(\theta+\varepsilon_s+a_s-a_s^*\left(y^{t-1}\right)\right)+a_t^*\left(y^{t-1}\right).\end{aligned}$$

So, a unit of effort in period s increases the agent's wage in each future period t by $\tau_\varepsilon/\tau_t=\tau_\varepsilon/\left(\tau_\theta+t\tau_\varepsilon\right)$.

Optimal effort is given by

$$c'\left(a_t\right)=(1-\delta)\sum_{s=t+1}^{\infty}\delta^{s-t}\frac{\tau_\varepsilon}{\tau_\theta+s\tau_\varepsilon}=:\gamma_t.$$

Implications

- γ_t is deterministic, independent of ε's.
  Agent works only to influence market's perception of ability.
  Due to normal signals, market uncertainty about agent ability is independent of ε's, hence so is effort.
- γ_t is decreasing, hence so is a_t. Market uncertainty about ability decreases over time.
- As $t \to \infty$, $\gamma_t \to 0$, hence so does a_t. Market uncertainty about ability eventually vanishes.
- γ_t depends on τ_θ and τ_ε only through their ratio $\tau_\theta / \tau_\varepsilon$, and is decreasing in this ratio.
  More uncertainty about $\theta \implies$ more to signal $\implies$ higher effort. More noise $\implies$ less gain from signaling $\implies$ lower effort.

Evolving theta

Suppose that instead of being perfectly persistent, ability is given by an AR(1) process $\theta_{t+1} = \theta_t + \beta_t$, with $\beta_t \sim N\left(0, \tau_\beta\right)$ iid.

Can show that any $\tau_\beta < \infty$ leads to a steady state with positive effort, where steady-state effort is a decreasing function of $\tau_\beta / \tau_\varepsilon$.

- SS effort is higher when there is more persistent uncertainty in ability, relative to noise.

Mailath-Samuelson 01

Consider an imperfect monitoring version of product choice game.

Similar to Holmström 82, but

- Agent knows θ, so more of a true reputation model.
- Assume binary type/effort/output, rather than continuous type/effort/output + normality as in Holmström.
- Binary usually more tractable when θ is private information.

Mailath-Samuelson Model

- Firm (agent) is either *normal* (prob μ) or *inept* (prob $1-\mu$).
- Firm's type is redrawn iid each period w/ prob λ. ("Renewal process.")
- Firm's type is its private information.
- Each period t, firm chooses effort $a \in \{0,1\}$ if normal, taking $a=1$ costs $c>0$. Inept firm always takes $a=0$.
- Effort stochastically determines the firm's period t *quality* y. Quality is high with prob p_H if $a=1$, with prob $p_L < p_H$ if $a=0$.
- Consumers (market) buy from firm at price $=\mathbb{E}\left[\text{quality}|h^t\right]$, where h^t is the history of past qualities.
- Firm has discount factor $\delta \in (0,1)$.

“Reputation as Separation”

Mailath-Samuelson emphasize that their model is one of “reputation as separation”: normal type wants to separate from inept type.

- They contrast this with FL’s “reputation as pooling”: rational type wants to pool with Stackelberg type.

However, this distinction breaks down in models with more than 2 types.

- In Holmström, agent works to appear higher-ability, i.e. both “separate” from low types and “pool” with high types.
- Diamond 89 is another early model with multiple types: his model has a commitment bad type, a commitment good type, and a normal type.

Equilibrium

Mailath-Samuelson ask whether there is a **Markov perfect eqm (MPE)** with high effort.

- Here, "Markov" means $\Pr(a_t = 1)$ depends only on firm's reputation, i.e. market's belief that $\theta = \mathit{normal}$.
- As in Holmström, there can also be non-Markov equilibria, which involve complicated mixing.

Theorem

For all $\lambda \in (0, 1)$ and $\mu > 0$, there exists $\bar{c} > 0$ s.t. if $c < \bar{c}$ there exists a MPE where the firm always exerts high effort.

- If market always expects $a_t = 1$, then observing y_L discretely lowers the firm's reputation μ_t (which is bounded strictly away from 0 and 1 due to replacements).
- Market price equals $\mu_t p_H$, strictly increasing in μ_t.
- If c is small enough, it's optimal for the firm to conform to the market's expectation of high effort.

Remarks

The eqm is not unique. It's also a MPE for the firm to always take $a = 0$, as then market price is 0 regardless of μ_t.

If $\lambda = 1$, the firm always takes $a = 0$ in every SE, as $\mu_t = \mu$ independent of history.

If $\lambda = 0$, the firm always takes $a = 0$ in every pure strategy SE (and hence in every MPE).

- Harder to show, but spirit is somewhat similar to Cripps-Mailath-Samuelson 04.
- Intuition: In a MPE, if firm takes $a = 1$ at reputation $\bar{\mu}$, also takes $a = 1$ any any $\mu' > \bar{\mu}$. (As observing y_L drops reputation/continuation value by more when it's higher.)
- But then after a long run of y_H's, firm's reputation will be very close to 1, and is therefore "entrenched." At this point, firm can safely deviate to $a = 0$, which breaks the eqm.

Remarks (cntd.)

“No effort without replacements” also relates to Holmström.

Both models have no long-run effort without replacements, but Holmström has short-run effort. Why?

In MS, if the normal firm exerts effort infinitely often, consumers eventually learn its type, so it no longer has an incentive to exert effort. So no long-run effort, as in Holmström.

Difference: in MS, a high type that doesn't exert effort is no more productive than a low type, so since high types eventually stop working, the value of a good reputation unravels, and there's also no short-run effort.

In Holmström, high types are always more productive than low types, so a good reputation is valuable even though high types eventually stop working. This incentivizes short-run effort.

Brand Names: Reputation as a Tradable Asset

Following Tadelis 99, Mailath-Samuelson also investigate "markets *for* reputation," where firms have brand names and randomly exit/enter, and can purchase existing brand names when they enter.

They find that new competant firms buy intermediate reputations and try to build them up, while new inept firms either buy low reputations or buy high reputations and run them down.

These results are important for understanding the dynamics of firm reputation and the value of brands. See the papers or MS Ch. 18.7 if curious.

Reputation for Quality (Board-Meyer-ter-Vehn 13)

Model a firm's reputation as the market's belief about an **endogenous type**: the firm's current quality, which is stochastically determined by past investments.

In Holmström or Mailath-Samuelson, my experience at your restaurant depends on whether you're a good cook (θ, exogenous) and whether you work hard tonight (a_t). In Board-Meyer-ter-Vehn, it depends on whether your restaurant is currently in a "high-quality state" (θ_t), which is stochastically determined by past investments ($(a_s)_{s<t}$).

Assume consumers get noisy signals of quality, and no signals of investments (other than via quality).

They find a tractable way of modeling this and analyzing the long-run dynamics of the firm's investment and reputation. The dynamics depend on the signal process by which consumers learn about quality.

Model

Time is continuous (more tractable). Firm has discount rate $r > 0$.

At each time t, firm has *quality* (type) $\theta_t \in \{L, H\}$, its private information. The firm gets a flow benefit (price) of $x_t = \mathbb{E}\left[\theta_t | h^t\right]$, where h^t is the public history of signals (see below).

At each time t, the firm chooses investment $a_t \in [0, 1]$ at flow cost ca_t (subject to measurability conditions).

Initial quality θ_0 is exogenous and arbitrary. At exogenous Poisson rate λ, a shock arrives (unobserved by consumers) which resets θ_t to H with prob a_t, L with prob $1 - a_t$.

- This is how effort matters: random quality-update times arrive, and the update is more likely to be to H if you happened to be investing at the arrival time.
  (Kind of strange, but tractable...)

Model (Cntd.)

Consumers learn about current quality θ_t through *another* Poisson process, this one with *endogenous* arrival rate μ_{θ_t}. There are two distinct cases:

1. *Good news signals:* $\mu_H > \mu_L$.
   Perfect good news is $\mu_L = 0$ (arrival proves $\theta_t = H$).
2. *Bad news signals:* $\mu_H < \mu_L$.
   Perfect bad news is $\mu_H = 0$ (arrival proves $\theta_t = L$).
3. (If $\mu_H = \mu_L$, signals are uninformative, and the unique eqm has $a_t = 0 \ \forall t$.)

Consider MPE where investment is $a_t = a(\theta_t, x_t)$. Expected investment according to market is
$\tilde{a}(x_t) = x_t a(H, x_t) + (1 - x_t) a(L, x_t)$.

Model summary: investment today stochastically determines future quality (via quality update shocks), which stochastically determines future reputation (via signal arrival rate μ_{θ_s}, $s > t$).

Equilibrium

Denote firm's value function by

$$V_\theta(x) = \int e^{-rt}\mathbb{E}^a\left[x_t - ca_t | \theta_0 = \theta\right] dt.$$

Let $\Delta(x) = V_H(x) - V_L(x)$.

- Value of actually having high-quality when reputation is x.

First result: $a(\theta, x)$ is independent of θ (since a only matters if quality updates today, in which case θ is obsolete) and given by

$$a(x) = \begin{cases} 1 & \text{if } \lambda\Delta(x) > c \\ 0 & \text{if } \lambda\Delta(x) < c \end{cases}.$$

Paper derives results on $\Delta(x)$ and dynamics for general case, but results sharpest for perfect good news and perfect bad news cases.

- Quite different between the two cases.

Perfect Good News

Arrival proves that $\theta_t = 1$, jumps x up to 1.

- Consumers learn through "breakthroughs."
- In absence of a breakthrough, x drifts down if $a(x) = 0$, ambiguous if $a(x) = 1$.

Breakthrough raises reputation to 1, so more valuable when current reputation is low. $\Delta(x)$ is decreasing in x.

Eqm is "work-shirk": $\exists$cutoff x^* s.t. $a(x) = 1$ if $x < x^*$, $a(x) = 0$ if $x > x^*$.

Reputation continually cycles: high-x firm shirks, reputation drifts down until $x \leq x^*$, then starts working, eventually reputation jumps up to 1.

- In work regime, depending on parameters, reputation can stabilize at x^* or drift down to a lower stationary point $x^g < x^*$.

For some parameters, there is a unique MPE.

Perfect Bad News

Arrival proves that $\theta_t = 0$, jumps x down to 0.

- Consumers learn through "breakdowns."
- In absence of a breakdown, x drifts up if $a(x) = 1$, ambiguous if $a(x) = 0$.

Breakdown drops reputation to 0, so avoiding breakdown more valuable when current reputation is high. $\Delta(x)$ is increasing in x.

Eqm is "shirk-work": $\exists$cutoff x^* s.t. $a(x) = 0$ if $x < x^*$, $a(x) = 1$ if $x > x^*$.

Reputation converges to 0 or 1: High-x firms work, never break down, x drifts up to 1. Low-x firms shirk, eventually break down (or start working if x drifts up to x^*), stop working forever if break down.

- In shirk regime, depending on parameters, x can drift up to x^* or drift down to a lower stationary point $x^b < x^*$.

For some parameters, there are a continuum of eqm thresholds x^*.

Bad Reputation: Background

In models considered so far, reputation effects increased efficiency by incentivizing greater effort.

But it's also possible for reputation effects to incentivize "misplaced effort," reducing efficiency.

Basic idea goes back at least to Scharfstein-Stein 90 on herding in financial markets.

- A financial advisor can be smart (gets accurate signal of which asset to buy) or not-so-smart (noisy signal).
- Advisor wants to convince the market that she's smart (career/reputation concern).
- If prior is that asset A is much likelier to be the right buy than B, then there's no eqm where the advisor always honestly reveals her signal.

  If there were, an advisor who recommends A is more likely to be smart than an advisor who recommends B.
- Without reputation concerns, honesty would be an eqm.

Morris 01

In Scharfstein-Stein, reputation concerns are exogenous: advisor is paid expected ability, so wants to look high-ability. (As in Holmström 82 and related papers like Brandenberger-Polak 96, Prendergast 93, Prendergast-Stole 96, Zwiebel 95.)

Morris 01 ("Political Correctness") endogenizes reputation concerns through a desire to be trusted in the future.

- State $\theta \in \{0, 1\}$, uninformed principal takes action $a \in \{0, 1\}$, wants to match state.
- Advsior knows θ can be good (preferences aligned with principal) or bad (always wants $a = 0$).
- Game is repeated twice, and both types of advisor care much more about the second period.
- Morris shows that the unique eqm of the first period is "babbling": both types of advisor choose the same distribution over recommendations in both states.
- Intuition: in a separating eqm, each type of advisor would make whichever recommendation improved her reputation.

Ely-Valimäki 03: Timing

EV 03 is a striking example of bad reputation. It considers a Morris-type model with an infinite time horizon and makes sharp negative predictions.

LR player 1 (firm, "mechanic") faces series of SR player 2's (consumers, "motorists").

Each consumer decides whether to hire the mechanic or exit and take outside option of 0.

If hires mechanic, mechanism observes if the car needs a **tuneup** ($\theta = \theta_t$) or a new **engine** ($\theta = \theta_e$), equally likely.

The mechanic then decides whether to do a tuneup ($a = t$) or replace the engine ($a = e$).

EV 03: Preferences + Information

The consumer gets $u > 0$ if mechanic takes the right action (a in state θ_a), gets $w < 0$ if mechanic takes the wrong action.

- Assume $w > u$, so consumer takes outside option if a and θ are uncorrelated.

Mechanic is *good* w/ prob $1 - \mu_0$, *bad* w/ prob μ_0.

- If good, stage game payoffs are identical to the consumer's, maximizes discounted utility with $\delta \in (0, 1)$.
- If bad, always takes $a = e$. (This is a *commitment bad type*. EV separately consider a *strategic bad type* with stage game payoff $\mathbf{1}\{a = e\}$, discount factor δ.)

Consumers observe the history of actions (hiring decisions and repairs) but not states or payoffs.

Intuition for Bad Reputation

If mechanic is thought to be bad (e.g. if takes e 100 times in a row), consumer won't hire.

Hence, if mechanic is good, at *pivotal histories* (e.g. if took e 99 times in a row) has an incentive to take t even in state θ_e to prove that he's good.

- Even though he doesn't like taking t in state θ_e.
- Key feature: good mechanic doesn't separate by exerting **more** effort (good for consumer, as in Holmström or Mailath-Samuelson), but by taking "the **wrong kind** of effort" (bad for consumer).

Intuition for Bad Reputation (cntd.)

Since the consumer is a SR player, at pivotal histories she won't hire the mechanic **even if she thinks he's good**, since she knows he'll always take t if she hires.

- Indeed, she won't hire at these histories for **any** belief, because her payoff when the mechanic is bad is also negative.

But then the eqm unravels, and the mechanic is **never** hired even if μ_0 is very small: if 99 e's is pivotal, consumer won't hire after 99 e's; but then 98 e's becomes pivotal. . .

Bad Reputation Theorem: Warm-Up Version

Instructive to first restrict attention to **renegotiation-proof eqm**, which here means that the mechanic is always hired at any history where he is known to be good.

Let $\delta^* = (u + w) / (2u + w)$. This is the cutoff δ^* s.t. if $\delta \geq \delta^*$ then $(1 - \delta)(-w) + \delta(u) \geq (1 - \delta)(u) + \delta(0)$, so mechanic will take the wrong action if this establishes his reputation.

Theorem

If $\delta > \delta^*$, then the mechanic is never hired in any renegotiation-proof NE.

- This theorem also holds for a strategic bad type.

Proof

- Fix a NE. Let $\bar{\mu}$ be the sup of the set of beliefs at which the mech is hired with positive prob. Suppose toward a contradiction that $\bar{\mu} > 0$.
- Whenever mech is hired with positive prob, he must take the right action in each state with prob bounded away from 0. (Otherwise, the consumer wouldn't hire.) So, by Bayes' rule, if he takes e at such a history, his reputation increases by a discrete amount.
- Hence, there exists a history where (1) μ is just below $\bar{\mu}$, (2) mech is hired with positive prob, (3) in state θ_e, mech takes e with positive prob, and doing so increases his reputation above $\bar{\mu}$. In this case, mech is never hired again (by defn of $\bar{\mu}$), so his payoff is $(1-\delta)(u)$.
- But if mech deviates to t, he proves he's good and hence is always hired in future (by renegotiation-proofness), so his payoff is $(1-\delta)(-w)+\delta(u)$.
- Since $\delta > \delta^*$, the deviation is strictly profitable.

Bad Reputation Theorem: Main Version

Without renegotiation-proofness refinement, can be NE where the mechanic is sometimes hired, but EV show that nonetheless payoffs for the consumers/good mechanic go to 0 as $\delta \to 1$.

Theorem

Let $\bar{V}(\mu, \delta)$ be the sup *of the good mechanic's payoff over all NE with commitment prob $\mu > 0$ and discount factor δ. Then $\lim_{\delta \to 1} \bar{V}(\mu, \delta) = 0$.*

- Consumer hires only if prob that good mech takes right action in state θ_t is above some cutoff indep of δ.
- So there is a fixed number of e's beyond which good mech stops getting hired.
- If taking t before this leads to a "reasonably large" increase in future average prob that mech gets hired, then when he's patient he strictly prefers to take t in both states after some smaller number of e's, a contradiction.
- So on average mech can be hired only very rarely even when known to be good.

When is Reputation Bad?

Ely-Fudenberg-Levine 08 ask how general bad reputation is.

They have several ways of getting at this. One of the most interesting is asking what happens in the EV 03 game if we add Stackelberg types who always take the right action, in addition to the bad types.

Intuitively, if the Stackelberg type is much more likely than the bad type then the bad reputation result goes away, and there are equilibibiria with high payoffs for the good mechanic/consumers.

This turns out to be true, but interestingly when the prob of both the Stackelberg type and the bad type are small, the prob of the Stackelberg type must be much (in the limit, infinitely) greater than that of the bad type. So EV 03 is "almost prior independent."

- We'll see similar "almost prior independence" results when we cover reputational bargaining next week.

Sketch of Result

First suppose there are only Stackelberg types and bad types (so no good types). Let $\eta = \Pr(\mathit{Stackelberg})$. Then consumers will hire iff

$$\eta u + (1-\eta)\left(-\frac{w-u}{2}\right) \geq 0,$$
$$\text{or } \eta \geq \frac{w-u}{u+w} =: \eta^*$$

Now suppose there are also good types.

- Good types are weakly worse for consumers than Stackelberg types, so mech is never hired if $\Pr(\mathit{bad}) > 1-\eta^*$.
- Good types are weakly better for consumers than bad types, so mech is always hired if $\Pr(\mathit{Stackelberg}) > \eta^*$.
- Question is what happens when $\Pr(\mathit{bad}) < 1-\eta^*$ and $\Pr(\mathit{Stackelberg}) < \eta^*$.

Sketch (cntd.)

Ely-Fudenberg-Levine show that if the prior lies below a curve that divides the remaining region, then as in EV 03 the good type's payoff converges to 0 as $\delta \to 1$ in any NE.

The notable result is that the curve is "vertical" when $\Pr(good) = 1$, so almost any perturbation of the complete information game is below the curve.

Intuition: Opportunity to form a reputation for being the Stackelberg type doesn't help with EV's problem, which is that at pivotal histories the good type has an incentive to take $a = t$ in state θ_e to separate from bad types. So Stackelberg types only change EV's result if they're present with substantial probability, not just as a perturbation.