Endogenous Experimentation in Organizations

Germán Gieczewski **Svetlana Kosterina**

Princeton University

October 18, 2021

Motivation

Organizations often face opportunities to experiment

The common wisdom:

- Experimentation should respond to information
  - if outcomes are bad, agents become more pessimistic
  - if enough negative evidence accumulates, stop experimenting
- When experimentation is collective:
  - incentive conflicts discourage experimentation
    - temptation to free-ride
    - fears that information will be misused by other agents

Motivation

Hence organizations should experiment the right (finite) amount of time, or less than that

However, history is littered with examples of organizations that stubbornly stuck with unsuccessful policies to the bitter end

- Political parties & movements
- Firms

We provide an explanation for obstinate behavior in organizations driven by self-selection of agents into the organization

Example: Theranos

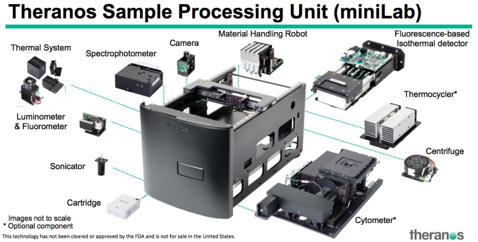

Theranos: startup founded in 2003 by Elizabeth Holmes

Goal: produce a portable machine that could run hundreds of lab tests on a single drop of blood

If successful, would revolutionize medical testing

Example: Theranos

Between its founding and 2013, Theranos stuck to a policy of pursuing Holmes's vision even as it seemed increasingly infeasible

> Theranos' chief scientist to his wife: "nothing at Theranos is working"

During this 10-year period most employees and board members became increasingly disillusioned

The board came close to removing Holmes as CEO

Example: Theranos

But it never did, because too many of the people who'd lost hope left instead

> Friend of a board member: "given everything you now know about this company, do you really want to own more of it?"

The final result: having developed no plan B if its moonshot didn't work, and not wanting to lose a profitable Walgreens partnership, Theranos launched in 2013 with innaccurate and fraudulent tests

Scandal destroyed the company

Example: European Communist parties

- Several countries in Europe had strong Communist parties in postwar period
- Their platform can be seen as a gamble on the state of the world
  - I claim: "capitalism is on the verge of collapse"
  - The state of the world: whether communism is fundamentally superior
  - If capitalism collapses, I am vindicated
  - If not, I will lose support
- The French Communist Party, for example, declined from roughly 20% electoral support in postwar period to less than 3% in late 2010s
- Dramatic decline, related to Soviet Union collapse, but didn't make the party abandon its core tenets

More examples

Premises

In this paper, we provide an explanation of obstinate behavior by organizations that rests on the following premises

Uncertain policy quality

organizations are uncertain about policy quality and must experiment to learn

Diverse beliefs

people often disagree about the merits of a policy

Changing membership

those who are more optimistic about an organization's policy self-select in

Endogenous policy

policy decisions are made by the current members of the organization

Results

Compared to a world in which the same agent always controls the policy, over-experimentation is possible

- we provide necessary and sufficient conditions under which the organization uses the risky policy forever in equilibrium, regardless of the outcome

Organization's policy may respond to information in a perverse way

- organization may be more likely to experiment forever if the risky policy is bad than if it is good

Main Forces

Two forces affect the amount of experimentation:

Changes in membership due to experimentation:

- when results are good, more agents come in (at the margin, pessimists)
- when results are bad, only optimists stay
- this dampens policy responsiveness to learning

Current pivotal member fears loss of control over future experimentation:

- side effect of changing membership
- leads to less experimentation

Related Literature

Exit, Voice and Loyalty (Hirschman, 1970)

Experimentation with multiple agents: Keller, Rady, and Cripps (2005), Keller and Rady (2010, 2015), Strulovici (2010)

- KRC (2005): underexperimentation due to free-riding (not in our model)
- Strulovici (2010): underexperimentation due to imperfect control over policy after learning (similar to our second effect)

Dynamic decision-making in clubs: Acemoglu, Egorov, and Sonin (2008, 2012, 2015), Bai and Lagunoff (2011), Gieczewski (2021)

- These are about heterogeneous preferences – no learning
- Because with het priors learning leads to agreement, we might expect that in the long run the policy pleases most agents – **not** the case
- Our main result, that there can be perpetual experimentation that pleases fewer and fewer people, only possible with learning

The Model

Time is continuous

Organization has a *risky* policy and a *safe* policy

The risky policy is either *good* or *bad*

Continuum of agents: continuous density f with support $[0, 1]$

- Agent $x \in [0, 1]$ has a prior belief that the risky policy is good with prob x

Agents discount future at rate γ

The Model

At every instant, each agent chooses whether to be a member of the organization

If an agent is not a member at t, she works independently and gets an *autarkic* flow payoff a

If she is a member, her payoff depends on the policy:

- under the safe policy: flow payoff s

- under the risky policy: zero flow payoff, except when the policy succeeds

The Model

successes: Poisson process with rate λ under the good risky policy

bad risky policy never succeeds

each time the risky policy succeeds, all members get a lump-sum payoff of size h

denote $g = \lambda h$

The Model

$0 < a < s < g$

safe policy is better than the outside option

good risky policy is the best

bad risky policy is the worst

The Model

Successes are observed by everyone

Agents decide whether to be members based on flow payoffs:

- If risky policy is being used, an agent with posterior p_t at time t wants to be a member if $p_t g \geq a$, so set of members is of the form $[y_t, 1]$, with $p_t(y_t) = \frac{a}{g}$
- If safe policy is being used, all agents want to be members ($s > a$)

The Model

At each t, the median member (induced by the current policy *from the immediate past*) chooses whether the organization experiments

> i.e., for experimentation to stop at time t, the median of the set $[y_t, 1]$ must want to stop – we rule out a mass of agents entering simultaneously and immediately changing the policy

Solve for MPE (state is information acquired from experimentation)

technical details

Some Observations

If the risky policy is used and no successes have been observed, then

- All agents become more pessimistic over time
- The set of members at t is of the form $[y_t, 1]$, with y_t increasing in t
- The policy choice at t is made by m_t, the median of $[y_t, 1]$

Some Observations

If a success is observed, then:

- the posterior of all agents jumps to 1
- all agents join and stay forever
- the risky policy is used forever

If the safe policy is ever used, then:

- all agents join
- the median becomes more pessimistic (pessimists join)
- no further learning happens
- the safe policy is used forever

So, assuming that the risky policy is used at $t = 0$, we just have to figure out when/if the organization switches to the safe policy, conditional on no successes

Equilibrium Characterization

Denote:

$V(x)$: continuation utility of an agent with posterior x at t provided that she expects experimentation to continue for all $t' \geq t$

m_t: median voter at t if the organization has experimented up to t

$p_t(m_t)$: m_t's posterior given no successes up to t

Note: $V(x)$, m_t, $p_t(m_t)$ are **not** equilibrium objects—they are calculated under a specific history and expected continuation

Equilibrium Characterization

Proposition

If $\inf_{t\geq 0} V(p_t(m_t)) > \frac{s}{\gamma}$*, there is a unique equilibrium. In it, the organization experiments forever.*

If $\inf_{t\geq 0} V(p_t(m_t)) < \frac{s}{\gamma}$*, the organization cannot experiment forever in equilibrium.*

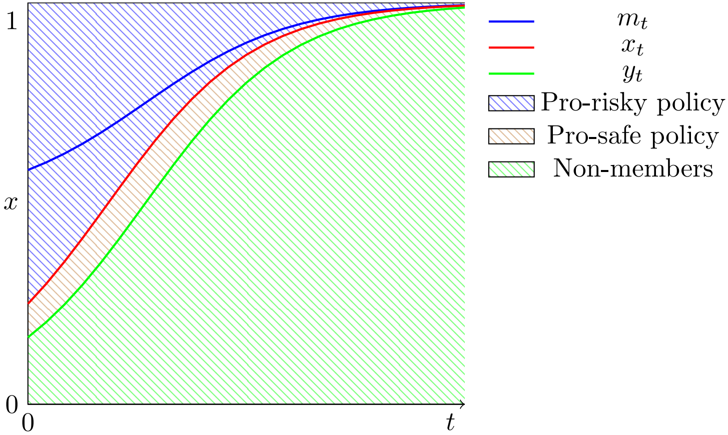

Figure 1: Median voter, indifferent voter, and marginal member on the equilibrium path

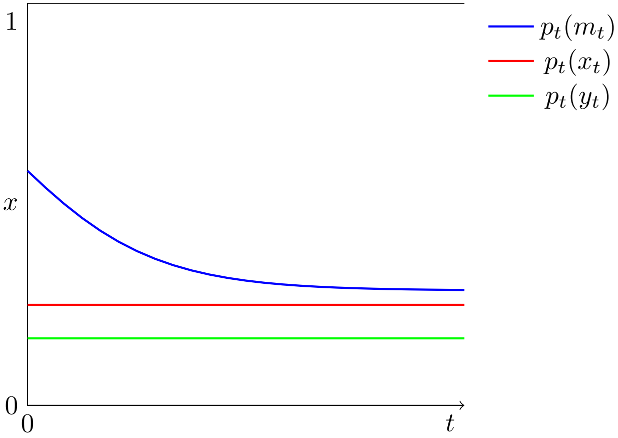

Figure 2: Posterior beliefs on the equilibrium path

Intuition: Why is perpetual experimentation possible?

As experimentation continues without success, posteriors of all agents tend to zero as $t \to \infty$

Pessimistic agents prefer safe policy, but agents with posteriors below $\frac{a}{g}$ leave, so they no longer vote

Hence median posterior does *not* converge to zero

If the org has experimented until t, and the pivotal agent, m_t, expects experimentation to continue forever if she does not stop it now, her choice is between $V(p_t(m_t))$ and $\frac{s}{\gamma}$

Hence, if $V(p_t(m_t)) > \frac{s}{\gamma}\ \forall t$, perpetual experimentation is an equilibrium

Why is perpetual experimentation the unique equilibrium?

Why are there no other equilibria?

Denote by $W_T(x)$ the payoff of an agent with posterior x if she expects experimentation for time T, then switch to safe policy if no successes

$W_0(x) \equiv \frac{s}{\gamma}$, $\lim_{T\to\infty} W_T(x) = V(x)$, and crucially, $W_T(x)$ can be shown to be single-peaked in T

Hence, if $V(p_t(m_t)) > \frac{s}{\gamma}$, m_t will not want to stop *no matter what continuation she expects*

Equilibrium Characterization

Proposition

The value function V in Proposition 1 satisfies the following:

(i) *If f is non-decreasing, then*

$$\gamma \inf_{t \geq 0} V(p_t(m_t)) = \gamma V\left(\frac{2a}{g+a}\right) = \frac{2ga}{g+a} + \left(\frac{1}{2}\right)^{\frac{\gamma}{\lambda}} \frac{a(g-a)}{g+a} \frac{\lambda}{\gamma+\lambda}.$$

Equilibrium Characterization

Let $f_\omega(x) = (\omega+1)(1-x)^\omega$ for $x \in [0,1]$, $\omega > 0$

Denote $\eta = \frac{1}{2^{\frac{1}{\omega+1}}}$

Proposition

(ii) *If the distribution of the agents' priors, f, MLRP-dominates f_ω, then*

$$\gamma \inf_{t\geq 0} V(p_t(m_t)) \geq$$

$$\gamma V\left(\frac{a}{\eta g+(1-\eta)a}\right) = \frac{ga}{\eta g+(1-\eta)a} + \eta^{\frac{\gamma+\lambda}{\lambda}} \frac{a(g-a)}{\eta g+(1-\eta)a}\frac{\lambda}{\gamma+\lambda}$$

with equality if $f = f_\omega$.

Equilibrium Characterization

Proposition

(iii) *Let f be any density with support $[0, 1]$. Then*

$$\gamma \inf_{t \geq 0} V(p_t(m_t)) \geq \gamma V\left(\frac{a}{g}\right) = a + \frac{a(g-a)}{g}\frac{\lambda}{\gamma+\lambda}$$

Corollary

Holding all other parameters constant, if a is close enough to s, the organization experiments forever.

(More precisely, this holds whenever $a > \frac{s}{1+\frac{g-s}{g}\frac{\lambda}{\gamma+\lambda}}$.)

Intuition

$$\gamma V\left(\frac{2a}{g+a}\right) = \frac{2ga}{g+a} + \left(\frac{1}{2}\right)^{\frac{\gamma}{\lambda}} \frac{a(g-a)}{g+a}\frac{\lambda}{\gamma+\lambda}$$

First term: median voter's expected flow payoff from risky policy, for large t

Second term: option value of experimenting

- Lower than if the current decision-maker had full control over future policy
- But still positive because the agent can quit if she becomes pessimistic enough (as long as $a > 0$)

Intuition

Sufficient condition for perpetual experimentation is not difficult to satisfy

More likely to hold when

- returns from good risky policy are high
- agents are patient
- outside option is close to safe policy

E.g., if $s = 3$, $a = 2$, $h = 1$ and f is uniform:

- when $g \geq 6$, it holds regardless of γ
- when $3.33 < g < 6$, it holds for low enough γ
- when $g \leq 3.33$, it cannot hold

What happens otherwise?

When perpetual experimentation is not possible, there are multiple equilibria

Both over-experimentation and under-experimentation are possible

Details

Other Learning Processes

So far, we have studied a model of *good news* that are perfectly informative

The main results extend to more general learning processes

- In fact, the potential for perverse effects expands

We show the case of imperfectly informative good news, but results are more general

- Different (and, in a sense, more special) results for perfectly informative bad news

Details

Imperfectly Informative Good News

Same model, only the risky policy generates different flow payoffs:

good risky policy: flow payoff 0, succeeds at rate λ, lump sums h

bad risky policy: flow payoff 0, succeeds at rate λ', with $b = \lambda' h$ and $b < a < s < g$

Posterior beliefs are interior even after a success, and a sufficient statistic for them (at time t, after k successes) is

$$L(k,t) = \left(\frac{\lambda'}{\lambda}\right)^k e^{(\lambda-\lambda')t}$$

We look for Markov equilibria, given by a set of values of L for which experimentation stops

Equilibrium: Imperfectly Informative Good News

Proposition

If $\inf_{L\geq 0} V(p_L(m_L)) > \frac{s}{\gamma}$*, there is a unique equilibrium. In it, the organization experiments forever.*

If $\inf_{L\geq 0} V(p_L(m_L)) < \frac{s}{\gamma}$*, the organization cannot experiment forever in equilibrium.*

We can give explicit bounds: $V(x) \geq \frac{xg+(1-x)b}{\gamma}$, and if f is non-decreasing, $\inf_{L\geq 0} p_L(m_L) \geq \frac{2(a-b)}{(g-b)+(a-b)}$ (extendable to other densities as before)

In particular, perpetual experimentation always obtains if s is close enough to a.

Nonmonotonic Equilibria

Proposition

For appropriately chosen g, s, a, v and f, $\epsilon \in (0, 1]$ and $L^* > 0$, an equilibrium of the following form exists:

whenever $L_t = L^*$, the organization stops experimenting with probability ϵ otherwise, the organization always continues experimenting.

When state L^* is below the initial state, experimentation can only stop if it's succeeded relatively often

Corollary: experimentation more likely to stop when the state of the world is good

Extensions: Heterogeneous payoffs

In some settings (e.g., if risky policy is a change in political institutions), payoffs may be heterogeneous ex post

> winners get payoffs, losers do not (Strulovici, 2010)

Formally, take the baseline model, but now instead of agent x having a prior belief x that the risky policy is good (for everyone), she has a probability x of being a winner

Assumption: voter types *and* successes are uncorrelated

As the organization experiments, some agents succeed $\Longrightarrow$ become sure winners, others remain unsure voters (progressively more pessimistic)

> sure winners never leave (& always want experimentation)
> unsure voters eventually leave

Heterogeneous payoffs

What is the effect of self-selection in this case?

To answer this, useful to first consider an intermediate step: baseline model **with no reentry**

Baseline model without reentry generates perpetual experimentation for a smaller set of parameter values, for two reasons

- Experimentation less valuable because once you exit, cannot reenter (less option value)
- Exit later because remaining a member has option value $\Longrightarrow$ pessimists stay for longer

and yet...

Proposition

Holding all other parameters constant, if a is close enough to s, perpetual experimentation is an equilibrium of the no-reentry model.

Heterogeneous payoffs

The case with heterogeneous payoffs differs from the baseline model in two ways

- If you leave, you don't learn about your type: equivalent to no reentry (discourages experimentation)
- Some agents become sure winners—as they are revealed, unshakable base of support for risky policy grows (supports experimentation)

This case generates perpetual experimentation in a larger set of parameter values than the no-reentry case, due to existence of sure winners

Comparison with baseline model is ambiguous

Extensions: Tradable shares

If the organization is a publicly-traded firm, and members are **investors**, members can own varying number of shares and trade them freely at equilibrium price

Suppose for now that the total number of shares (& org size) is **fixed**

> (firm does not adjust share offerings/buybacks/dividends based on share demand)

Question: can selection (by buying more shares) cause over-experimentation?

> **Yes—but not infinite experimentation**

Tradable shares

Assume risk-averse agents: concave flow utility from consumption $u(c)$

- Otherwise, agent with belief 1 buys the firm
- For most results, assume CRRA utility: $u(c) = \frac{c^{1-\theta}}{1-\theta}$
- CARA utility also somewhat tractable

Add a credit market to the model: can lend & borrow at rate γ (=discount rate)

Also add an infinite mass of skeptical agents: agents with prior 0 (but belief still jumps to 1 after success)

- Simplifies analysis after success: constant price $\frac{g}{\gamma}$

Tradable shares

Unit mass of agents with density f as before

- $c_t(x)$: x's consumption at time t, if no success
- $c_t(x, \text{succ})$: x's consumption after a time-t success, constant
- $q_t(x)$: x's share demand at time t
- W_0: initial wealth for all agents
- $W_t(x)$: x's wealth at time t, including value of shares

In addition to individual decisions, m_t, the share-weighted median member at t, chooses whether to stop experimenting at each t

Tradable shares

Unit mass of shares

Market-clearing condition: $\int_0^1 f(x) q_t(x) = 1$

Eq price ρ_t if no success by time t, $\overline{\rho} = \frac{g}{\gamma}$ after success, $\frac{s}{\gamma}$ if safe policy

If x owns $q_t(x)$ shares, gets $q_t(x)h$ payout per success, $q_t(x)s$ flow payoff if safe policy being used

$h + \overline{\rho} - \rho_t$: net gain from success (capital gains/losses matter too!)

$Q_t(x) = q_t(x)(h + \overline{\rho} - \rho_t)$ normalized share demand at time t

Lottery tickets

$\xi_t = \frac{\gamma \rho_t - \rho_t'}{h + \overline{\rho} - \rho_t}$ flow cost of holding lottery tickets

Note: "outside option" is simply owning no/fewer shares

Tradable shares

Proposition

Assume CRRA utility with parameter θ. Then

(i) *There is no equilibrium with perpetual experimentation.*

(ii) $q_t(x)$ *is increasing in* x *for all* t.

(iii) *Moreover, if* $\theta = 1$, $q_t(x)$ *is MLRP-increasing in* t, x. *In particular,* m_t *is increasing in* t.

Intuition

Selection still at play: $q_t(x)$ increasing in x (ii)

Even with fixed number of shares, selection gets **stronger** over time (iii) (cf. baseline model)

Why?

Optimists' beliefs change more slowly: if $x' > x$, then $\frac{p_t(x')}{p_t(x)}$ increasing in t

Intuition

But this selection effect **not** strong enough to support perpetual experimentation (i)

- To incentivize agents with high enough posteriors to hold a majority of shares, price would have to be so low that less-optimistic agents would also demand shares—too many
- Selection mitigated by consumption smoothing: optimists expect to get rich soon $\Longrightarrow$ borrow and spend wealth—if no success, become poor, which lowers share demands
- Even with agent turnover, selection not strong enough for perpetual exp
- However, exponential population growth could support it
  - startups
- Or if current members could hand-pick new owners (possible in privately held firms)

Proof sketch

Let $V_t(W,x)$ be x's value function if no success, $U_t(W,x)$ after success, if current wealth is W

FOCs:

$$
\begin{aligned}
0 =& \gamma u'(c_t(x)) - \frac{\partial V_t(W_t(x),x)}{\partial W} \\
0 =& \lambda p_t(x)(h+\overline{\rho}-\rho_t)\frac{\partial U_t(W_t(x)+q_t(x)(h+\overline{\rho}-\rho_t),x)}{\partial W} + \\
& + \frac{\partial V_t(W_t(x),x)}{\partial W}(\rho_t' - \gamma\rho_t) \\
-\frac{\partial u'(c_t(x))}{\partial t} =& \lambda p_t(x)(u'(c_t(x,\text{succ})) - u'(c_t(x)))
\end{aligned}
$$

Proof sketch

Plugging the first FOC into the second, we get

$$c_t(x)^{-\theta}(\gamma\rho_t - \rho_t') = \lambda p_t(x)(h + \overline{\rho} - \rho_t)c_t(x;\text{succ})^{-\theta}$$

$$\left[\frac{c_t(x,\text{succ})}{c_t(x)}\right]^{\theta} = \frac{\lambda p_t(x)(h + \overline{\rho} - \rho_t)}{\gamma\rho_t - \rho_t'}$$

$$c_t(x,\text{succ}) = c_t(x)\left[\frac{\lambda p_t(x)}{\xi_t}\right]^{\frac{1}{\theta}}$$

Proof sketch

Plugging in the functional form for u' into the third FOC,

$$
\begin{aligned}
-\frac{\partial u'(c_t(x))}{\partial t} &= \lambda p_t(x)(c_t(x, succ)^{-\theta} - c_t(x)^{-\theta}) \\
-\frac{\partial u'(c_t(x))}{\partial t} &= \lambda p_t(x) c_t(x)^{-\theta} \left(\frac{\xi_t}{\lambda p_t(x)} - 1 \right) \\
-\frac{\frac{\partial c_t(x)^{-\theta}}{\partial t}}{c_t(x)^{-\theta}} &= \lambda p_t(x) \left(\frac{\xi_t}{\lambda p_t(x)} - 1 \right) \\
\theta \hat{c}_t(x) := \theta \frac{\frac{\partial c_t(x)}{\partial t}}{c_t(x)} &= \xi_t - \lambda p_t(x)
\end{aligned}
$$

where $\hat{f} = \frac{f'}{f}$

Proof sketch

Using that $c_t(x, succ) = c_t(x) \left[\frac{\lambda p_t(x)}{\xi_t} \right]^{\frac{1}{\theta}}$, we get

$$\hat{c}_t(x, succ) = \hat{c}_t(x) + \frac{1}{\theta} \left[\hat{p}_t(x) - \hat{\xi}_t \right]$$
$$\hat{c}_t(x, succ) = \frac{1}{\theta} \left[\xi_t - \lambda p_t(x) \right] + \frac{1}{\theta} \left[-\lambda(1 - p_t(x)) - \hat{\xi}_t \right]$$
$$\hat{c}_t(x, succ) = \frac{1}{\theta} \left[-\lambda + \xi_t - \hat{\xi}_t \right] =: \Gamma_t$$

Summarizing,

$$\theta \hat{c}_t(x) = \xi_t - \lambda p_t(x)$$
$$\theta \hat{c}_t(x, succ) = -\lambda + \xi_t - \hat{\xi}_t = \theta \Gamma_t$$

Proof sketch

Using the previous results plus the budget constraints:

$$W'_t(x) = \gamma W_t(x) - c_t(x) - Q_t(x)\xi_t$$
$$c_t(x, succ) = \gamma W_t(x) + \gamma Q_t(x)$$

we can express the entire path of x's behavior as a function of the price path $(\xi_t)_t$ and when x stops holding risky shares, $t^*(x)$

- either when $\lambda p_t(x) = \xi_t$
- or at t^*, when the organization switches to the safe policy—whichever is earlier

Proof sketch

Let $\zeta_t = \int_0^t \xi_s ds$. Denote $\psi_t = -\frac{\lambda t}{\theta} + \frac{\zeta_t}{\theta} - \gamma t - \zeta_t$.

$$c_t(x, succ) = \frac{W_0 e^{-\frac{\lambda t}{\theta} + \frac{\zeta_t}{\theta}} \left(\frac{1}{\xi_t}\right)^{\frac{1}{\theta}}}{\int_0^{t^*} e^{\psi_z} \left(\frac{1}{\xi_z}\right)^{\frac{1}{\theta}} \left[\left(\frac{\xi_z}{\lambda p_z(x)}\right)^{\frac{1}{\theta}} + \frac{\xi_z}{\gamma}\right] dz + \frac{1}{\gamma} e^{\psi_{t^*}} \left(\frac{1}{\lambda p_{t^*}(x)}\right)^{\frac{1}{\theta}}}$$

$$c_t(x) = \frac{W_0 e^{-\frac{\lambda t}{\theta} + \frac{\zeta_t}{\theta}} \left(\frac{1}{\lambda p_t(x)}\right)^{\frac{1}{\theta}}}{\int_0^{t^*} e^{\psi_z} \left(\frac{1}{\xi_z}\right)^{\frac{1}{\theta}} \left[\left(\frac{\xi_z}{\lambda p_z(x)}\right)^{\frac{1}{\theta}} + \frac{\xi_z}{\gamma}\right] dz + \frac{1}{\gamma} e^{\psi_{t^*}} \left(\frac{1}{\lambda p_{t^*}(x)}\right)^{\frac{1}{\theta}}}$$

$$W_t(x) = W_0 e^{\gamma t + \zeta_t} \left[1 - \frac{\int_0^t e^{\psi_z} \left(\frac{1}{\xi_z}\right)^{\frac{1}{\theta}} \left[\left(\frac{\xi_z}{\lambda p_z(x)}\right)^{\frac{1}{\theta}} + \frac{\xi_z}{\gamma}\right] dz}{\int_0^{t^*} e^{\psi_z} \left(\frac{1}{\xi_z}\right)^{\frac{1}{\theta}} \left[\left(\frac{\xi_z}{\lambda p_z(x)}\right)^{\frac{1}{\theta}} + \frac{\xi_z}{\gamma}\right] dz + \frac{1}{\gamma} e^{\psi_{t^*}} \left(\frac{1}{\lambda p_{t^*}(x)}\right)^{\frac{1}{\theta}}}\right]$$

$$Q_t(x) = W_0 e^{\gamma t + \zeta_t} \frac{\frac{1}{\gamma} e^{\psi_t} \left(\frac{1}{\xi_t}\right)^{\frac{1}{\theta}} - \int_t^{t^*} e^{\psi_z} \left(\frac{1}{\xi_z}\right)^{\frac{1}{\theta}} \left[\left(\frac{\xi_z}{\lambda p_z(x)}\right)^{\frac{1}{\theta}} + \frac{\xi_z}{\gamma}\right] dz - \frac{1}{\gamma} e^{\psi_{t^*}} \left(\frac{1}{\lambda p_{t^*}(x)}\right)^{\frac{1}{\theta}}}{\int_0^{t^*} e^{\psi_z} \left(\frac{1}{\xi_z}\right)^{\frac{1}{\theta}} \left[\left(\frac{\xi_z}{\lambda p_z(x)}\right)^{\frac{1}{\theta}} + \frac{\xi_z}{\gamma}\right] dz + \frac{1}{\gamma} e^{\psi_{t^*}} \left(\frac{1}{\lambda p_{t^*}(x)}\right)^{\frac{1}{\theta}}}$$

Proof sketch

Part (ii): from previous equation

Part (i): suppose an equilibrium with perpetual experimentation

- If price decreases faster than λ (formally, if there is a sequence $t_n \to \infty$ with $\frac{\xi_{t_n}}{e^{-\lambda t_n}} \to 0$) then we can show aggregate share demands go to ∞—contradiction
- If not, we can show from previous equations that share demands for optimists cannot increase at rate λ forever—in fact, will be bounded
- Hence $p_t(m_t) \to 0$
- Then check that m_t would prefer to stop experimenting for large t

Proof sketch

Part (iii): after **many** simplifications, get

$$Q_t(x) = \max\left\{W_0 e^{\zeta_t}\left[x\left(\frac{\lambda}{\xi_t}e^{-\lambda t} - e^{-\lambda t} + 1\right) - 1\right], 0\right\}$$

This is MLRP-increasing in t, x iff $A(t) = \frac{\lambda}{\xi_t}e^{-\lambda t} - e^{-\lambda t} + 1$ is decreasing in t

The market-clearing constraint is

$$W_0 e^{\zeta_t}\int_0^1 \max\{xA(t) - 1, 0\} f(x)dx = h + \overline{\rho} - \rho_t$$

The log-derivative of e^{ζ_t} is $\xi_t = \frac{\gamma\rho_t - \rho_t'}{h+\overline{\rho}-\rho_t}$, while the log-derivative of the RHS is $\frac{-\rho_t'}{h+\overline{\rho}-\rho_t}$, which is lower

hence $A(t)$ decreasing in t

Flexible size?

With fixed population, is perpetual experimentation again possible if organization size is flexible?

optimists could shrink the org as no success arrives

makes it easier for them to retain control?

No!

Flexible size?

Formally, assume: org has *scale of operations* $M_t \leq 1$ at time t if risky policy with no success by then

- means a success only produces a total payoff hM_t
- after success or switch to safe policy, scale returns to 1
- note: in baseline model, M_t is an equilibrium object, and the fact that it shrinks over time enables over-experimentation
- here we allow M_t to be a parameter

Corollary

?? *For any path $(M_t)_t$, there is no equilibrium with perpetual experimentation.*

Flexible size?

Key: **rights over future profits don't vanish**

Per-share gain from first success is $hM_t + \overline{\rho} - \rho_t$

- Even if org shrinks to negligible scale, value of success per share must always be at least $\overline{\rho} - \rho_t$
- Only h can be shrunken away, not capital gains: after success, org would grow back to full size

Capital markets may have a corrective effect on firms captured by optimists

Extensions

Other decision-making rules: analogous

Size-dependent payoffs

- Economies of scale, congestion effect: can go either way
- Minimum size (e.g. bankruptcy at time T): rules out perpetual exp, but still get exp up to time T
- Size-dependent learning rate: lowers V, but main results still hold

Two risky policies (instead of one risky & one safe)

- May still see perpetual use of first policy that is tried, or switching back and forth—still, inefficiently long spells

No free entry (think employees)

- Instead, assume size of organization is fixed; agents differ in prior beliefs and ability (uncorrelated); organization selects on ability
- Same results – but instead of shrinking, org selects worse agents

Thoughts

When s close to a, perpetual experimentation likely (but if $s = a$, not a problem)

When $s >> a$, enough pessimists stick around to change the policy

- The greater the availability of exit, the less likely voice will be used (Hirschman, 1970)

Model assumes away CEO agency

- Appropriate if strong control by investors/members/“the base”
- cf. Curley effect (Glaeser and Shleifer 2005)

Thank you!

Example: Blockbuster

Massive video rental chain

- $6B annual revenue, 84,000 employees at peak

Went bankrupt in 2010 amid competition from Netflix, Redbox

Failed to adapt to digital market—why?

Example: Blockbuster

In the early 2000s, Blockbuster management was optimistic about its business model and skeptical of DVD-by-mail and online streaming

CEO John Antioco: "I didn't believe that technology would threaten the company as fast as critics thought"

- This belief was one of his reasons for joining—already some self-selection

By 2004, Blockbuster's majority owner, Viacom, had lost faith in the company and sold its stake

Large share was bought by an activist investor skeptical of efforts to expand into the online market

Example: Blockbuster

By this point Antioco himself had decided Blockbuster had to compete in the online market, but now faced opposition from his board

In 2007, he launched a major expansion into the online market, and attempted a merger with Netflix

The board fired him and replaced him with a CEO who ended the online expansion and focused on brick-and-mortar stores

New CEO, in 2008:

> "Should we put shareholder money at risk in a market that's at best five years away from being commercial? I don't think so. [...] Neither Redbox nor Netflix are even on the radar screen in terms of competition."

Blockbuster went bankrupt 2 years later

Example: Communist regimes

Communist regimes themselves can be seen as a risky collective experiment (probably with heterogeneous payoffs)

Anecdotally, regimes that allowed or failed to prevent exit (Cuba, China) were more stable than those with stricter emigration restrictions (USSR & Eastern Bloc)

Emigrants appear to be self-selected: as many as 60% of Cuban Americans identify as Republican, compared to less than 30% for other Latinos

Return

Equilibrium Definition

If k successes after experimenting for time t, $L = L(k,t) = \mathbb{1}_{k=0}e^{-\lambda t}$

Strategies:

Membership function: $\beta(x, L, \pi) = 1$ if x is a member given information L and current policy π, 0 otherwise

Policy function: $\alpha(L, \pi) \subseteq \{0, 1\}$ is the set of median $m(L, \pi)$'s acceptable policies in state (L, π)

Analogous to a choice function in decision theory

Strategy Profile

Some notation:

$\left(\tilde{k}_\tau\right)_\tau$: Poisson process with rate b or 0 if $\theta = G$ or B respectively

$\tilde{L}_\tau = L\left(\tilde{k}_\tau, \tau\right)$

$n(t) = \int_0^t \pi_s ds$: the amount of experimentation up to time t

$\pi_{t_0^-}$, $\pi_{t_0^+}$ are the left- and right-limits of π_t at t_0

Strategy Profile

Strategy profile: membership function β, policy function α, and stochastic path of play consisting of information and policy paths $(L_t, \pi_t)_t$ s.t.:

(a) Given policy type θ, $(L_t, \pi_t)_{t \geq 0}$ is a progressively measurable Markov process with paths that have left and right limits at every $t \geq 0$, and such that $(L_0, \pi_0) = (1, 1)$.

(b) $L_t = \tilde{L}_{n(t)}$.

(c) $\pi_t \in \alpha(L_t, \pi_{t^-})$ for all $t \geq 0$.

(d) $\pi_{t^+} \in \alpha(L_t, \pi_t)$ for all $t \geq 0$.

Equilibrium Definition

Some notation:

$m(L, \pi)$: the median member of the organization given that information is L and the incumbent policy is π

$p(L, x)$: posterior of agent with prior x given L

$V_x(L, \pi)$: continuation utility of an agent with prior x given L and π

$\overline{V}_x(L, \pi, \epsilon)$: x's continuation utility starting from t with information L, when π is played during $[t, t + \epsilon)$ irrespective of the equilibrium strategies and the equilibrium strategies are played thereafter

Equilibrium Definition

An equilibrium is a strategy profile s.t.

(i) $\beta(x, L, \pi) = 1$ if $\pi p(L, x) g + (1 - \pi) s > a$ and 0 if $<$.

(ii) If $V_{m(L,\pi)}(L, \pi') > V_{m(L,\pi)}(L, 1 - \pi')$, then $\alpha(L, \pi) = \pi'$.

(iii) If $V_{m(L,\pi)}(L, 1) = V_{m(L,\pi)}(L, 0)$ but $\overline{V}_{m(L,\pi)}(L, \pi', \epsilon) - \overline{V}_{m(L,\pi)}(L, 1 - \pi', \epsilon) > 0$ for all $\epsilon > 0$ small enough, then $\alpha(L, \pi) = \pi'$.

Intuitive Version of the Equilibrium Definition

$m(L, \pi)$: median member given information L and incumbent policy π

Definition

An equilibrium satisfies the following:

(i) Agents choose whether to be members based on their flow payoffs.

(ii) If $m(L, \pi)$ prefers policy π' to the alternative, then policy π' is chosen.

(iii) If $m(L, \pi)$ gets the same continuation utility under either policy choice, but would get strictly higher utility from 'locking in' policy π' for any small length of time $\epsilon > 0$, then π' is chosen.

This is MPE: conditions only on current information L and current policy π (which determines the median voter).

Equilibrium Restriction

Consider a discrete-time game in which membership and policy decisions are made at $t \in \{0, \epsilon, 2\epsilon, \dots\}$ with $\epsilon > 0$ small

- At $t \in \{0, \epsilon, 2\epsilon, \dots\}$, first the median chooses a policy π_t for time $[t, t + \epsilon)$, then agents choose whether to be members
- Agents who choose to be members at t get flow payoffs from π_t during $[t, t + \epsilon)$ and are the incumbent members at $t + \epsilon$
- The median of this set chooses $\pi_{t+\epsilon}$

Small delay between joining the organization and voting on the policy rules out equilibria with self-fulfilling prophecies

Interested in equilibria that are limits of the equilibria of this game as $\epsilon \to 0$

slides

What happens if perpetual exp. is not an equilibrium?

Stopping function $\tau : [0, \infty) \to [0, \infty]$

$\tau(t) \geq t$ is s.t. m_t is indifferent about switching to the safe policy at t if she expects that, if she *doesn't* stop, experimentation will stop at $\tau(t)$

If she would always prefer to experiment, then $\tau(t) = \infty$; if strictly does not want to experiment, $\tau(t) = t$

$\tau(t)$ is unique

Equilibrium Characterization

Proposition

For some $t_0 \in [0, \tau(0)]$*, there is an equilibrium in which experimentation stops at* t_0*.*

If τ *is increasing, such an equilibrium exists for* ***every*** $t_0 \in [0, \tau(0)]$*.*

If $\tau(t) = \infty$ *for all* $t \in [0, T]$*, then experimentation continues up to at least time* T *in every equilibrium.*

Intuition

When agents are not willing to experiment forever, they are usually still willing to experiment for limited periods of time

Assume $t_0 < t_1 < \ldots$ are the times for which experimentation stops (if it has not stopped yet)

m_{t_n} must be indifferent between stopping and experimenting until t_{n+1}

Hence $t_{n+1} = \tau(t_n)$

m_0's optimal experimentation time is between 0 and $\tau(0)$, so from her point of view both over- and under-experimentation are possible

return

A Model of Bad News

Same model, only the risky policy generates different flow payoffs:

Good risky policy: constant flow payoff g

Bad risky policy: constant flow payoff g but occasionally experiences failures

Failures happen at rate λ and cost all members a lump sum of h

Expected flow payoffs are the same as before, but the learning process is different

Single-Agent Bandit Problem

Analogous single-agent bandit problem:

Good news: agent experiments up to some t, becomes more pessimistic, and stops then if no successes occur

Bad news: agent becomes more optimistic over time, and so experiments forever until a failure is observed (unless she would rather not start at all)

Let $\underline{t}$ be the earliest time when an agent with $V(p_t(m_t)) < \frac{s}{\gamma}$ is pivotal, if any such agent exists.

Equilibrium: A Model of Bad News

Proposition

1. *If* $V(p_t(m_t)) > \frac{s}{\gamma}$ *for all* t*, then there is a unique equilibrium. In it, the organization experiments forever.*

2. *If* $V(p_t(m_t)) < \frac{s}{\gamma}$ *for some* t*, then in any equilibrium the organization stops experimenting at a finite time* $T < \underline{t}$*.*

Intuition: A Model of Bad News

Experimenting forever is still possible

If not, generally there is under-experimentation, because here future and present experimentation are strategic complements

- If future agents do not experiment, this destroys some of my option value from experimentation
- Whenever an agent would choose not to experiment forever in the single-agent bandit problem, she would rather stop immediately
- And she can stop immediately in our model

Return