Dynamic Incentives with Brownian Information

Gonzalo Cisternas
MIT Sloan

Today

- Market-based incentives
  - Holmström (1999) & Cisternas (2017): learning, hidden actions and no commitment ("implicit incentives")
- Optimal incentive schemes
  - Sannikov (2008): hidden actions and commitment ("explicit incentives")

Frequent arrival of information

Market-Based Incentives

Holmström (1999) Random Walk/Stationary Learning

- Risk-neutral worker, competitive labor market, no explicit contracts

Discrete time	Continuous time
$\xi_t = \theta_t + a_t + \epsilon_t^\xi,\ \epsilon_t^\xi \sim \mathcal{N}(0,\sigma_\xi^2)$ $\theta_{t+1} = \theta_t + \epsilon_t^\theta,\ \epsilon_\theta \sim \mathcal{N}(0,\sigma_\theta^2)$	$d\xi_t = (a_t + \theta_t)dt + \sigma_\xi dZ_t^\xi$ $d\theta_t = \sigma_\theta dZ_t^\theta$, $Z^\xi \perp Z^\xi$ B.M.
$w_t := \mathbb{E}[\xi_t \| \xi^{t-1}] = \underbrace{\mathbb{E}[\theta_t \| \xi^{t-1}]}_{p_t^*} + a_t^*$	$w_t := \frac{\mathbb{E}[d\xi_t \| \xi^t]}{dt} = \underbrace{\mathbb{E}[\theta_t \| \xi^t]}_{p_t^*} + a_t^*$
$p_{t+1}^* = p_t^* + \frac{\gamma_t}{\gamma_t + \sigma_\xi^2}[\xi_t - (a_t^* + p_t^*)]$	$dp_t^* = \frac{\gamma_t}{\sigma_\xi^2}[d\xi_t - (a_t^* + p_t^*)dt]$
where $\gamma_t := \mathbb{E}[(\theta_t - p_t^*) \| \xi^{t-1}]$ ("$= 1/h_t$")	where $\gamma_t := \mathbb{E}[(\theta_t - p_t^*) \| \xi^t]$
$\gamma_t = \gamma_{t-1} + \sigma_\theta^2 - \frac{\gamma_t^2}{\gamma_t + \sigma_\xi^2}$ (check (14) & (15) p.173)	$\dot{\gamma}_t = \sigma_\theta^2 - \left(\frac{\gamma_t}{\sigma_\xi}\right)^2$
Payoff: $\mathbb{E}\Big[\sum_{t=0}^{\infty} \delta^t [w_t - g(a_t)]\Big]$	Payoff: $\mathbb{E}\Big[\int_{t=0}^{\infty} e^{-rt}[w_t - g(a_t)]dt\Big]$

- Now, focus on cts. time, with $\gamma_t = \gamma^* = \sigma_\theta^2 \sigma_\xi^2$ (stationary learning), letting $\beta := \gamma/\sigma_\xi^2$, so $dp_t^* = \beta[d\xi_t - (a_t^* + p_t^*)dt]$. Suppose a_t^* deterministic.

Holmström (1999) Random Walk/Stationary Learning

- $V_t := \mathbb{E}_t \left[\int_t^\infty e^{-(r-s)} [\underbrace{p_s^* + a_s^*}_{w_s =} - g(a_s)] ds \right]$, $t \geq 0$, *continuation value*
- Wage is linear in p_t^*. Moreover,

$$p_s^* = e^{-\beta(s-t)} p_t^* + \beta \underbrace{\int_t^s e^{-\beta(s-u)} d\xi_u}_{\text{linear in output history}} \text{ and } d\xi_t = \underbrace{(a_t + \theta_t)}_{\text{linear in effort}} dt + \sigma_\xi dZ_t^\xi$$

- Fully linear model $\Rightarrow$ pointwise optimization:

$$\begin{aligned} V_t &= \mathbb{E}_t \left[\int_t^\infty e^{-(r-s)} \left(\beta \int_t^s e^{-\beta(s-u)} a_u du - g(a_s) \right) ds \right] + \text{exog.} \\ &= \mathbb{E}_t \left[\int_t^\infty e^{-(r-s)} \left(\frac{\beta}{r+\beta} a_s - g(a_s) \right) ds \right] + \text{exog.} \end{aligned}$$

$$\Rightarrow \underbrace{g'(a^*)}_{\text{Mg. Cost}} = \underbrace{\frac{\beta}{\beta + r}}_{\text{Mg. Benefit}} = \underbrace{\beta}_{\frac{dp_t^*}{d\xi_t}} \times \underbrace{\frac{1}{r+\beta}}_{\substack{\text{total value of } \Delta p_t^*; \\ \text{measure of value of reputation}}}$$

Holmström (1999): Two Observations

1. Wage process $w_t = p_t^* + a^*$ is a function of the output history ξ^t via

$$p_t^*[\xi^t] = e^{-\beta t} p_0^* + \beta \int_0^t e^{-\beta(t-s)} d\xi_s$$

$\Rightarrow$ competition leads to an output-contingent contract

2. Such history-dependence is summarized in the value that the belief takes. In equilibrium

$$V_t^* = \frac{p_t^* + a^*}{r} + \frac{g(a^*)}{r} =: V^*(p_t^*)$$

But

$$\underbrace{V^{*\prime}(p_t^*) = \frac{1}{r}}_{\text{eqbm. mg. utility}} \neq \underbrace{\frac{1}{\beta + r} = q}_{\text{value of reputation}}$$

i.e. eqbm. marginal utility alone does not pin down behavior

Private Information and Incentives

- In equilibrium, p_t^* coincides with the worker's belief p_t
  - $\Rightarrow V^{*\prime}(\cdot)$ is really $\frac{d}{dp}[V(p,p^*)|_{p^*=p}]$
- $q = 1/(\beta + r)$ found by Holmström—a measure of marginal utility—cannot entail $p_t = p_t^*$ for all t
  - $p_t \neq p_t^* \Rightarrow p_t$ is a private belief that matters for incentives
- Homström's analysis is mathematically correct, but it does not address the issue of how private information can affect incentives
  - A private belief is never introduced (not necessary to find a^*)
  - a^* does not depend on beliefs (which can be misleading into suggesting that having pvte info does not affect the worker's behavior)
- Consequence: important economics missing

Cisternas (2017)

- "Two-Sided Learning and the Ratchet Principle," *Restud*, forthcoming

Holmström (1999)	Cisternas (2017)
market sets wage $w_t = p_t^* + a_t^*$	market's action $\chi(p_t^*, a_t^*)$

- Behavior affected by ratcheting forces: good performance today $\rightarrow$ market's expectation of future performance $\uparrow \Rightarrow$ incentives today $\downarrow$
  - Understand how private info. shapes the eqbm. found by Holmström
  - Beyond linear settings, additional effects affecting behavior
  - Expand the set of environments involving learning and unobserved actions that can be analyzed

- Focus on $\chi(p_t^*)$ today

Cisternas (2017)

- Recall that in Holmström $w_t^* = p_t^* + a_t^*$, so $dw_t^*/dp_t^* = 1$ (constant)
  - local benefit from reputation is constant $\Rightarrow a_t^* = a^*$

- Now, $\chi'(\cdot) \neq$ constant $\Rightarrow a_t^* = a^*(p_t^*)$ (Markov eqbm)

Holmström (1999)	Cisternas (2017)
$g'(a^*) = \beta q,\ q = \frac{1}{r+\beta}$	$g'(a^*(p_t^*)) = \beta q(p_t^*),\ p \mapsto q(p)?$

- But what is $q(\cdot)$ economically?

$q(\cdot)$: Value of a small degree of belief asymmetry

- Recall $V_t := \mathbb{E}_t\left[\int_t^\infty e^{-(r-s)}[\chi(p_t^*) - g(a_s)]ds\right]$
  - But if the market's conjecture is Markov, i.e., $a^*(p^*)$, $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot|(p_t^*, p_t)]$
  - Write $V_t = V(p_t^*, p_t)$ after any (private) history (a^t, ξ^t)

- Agent controls p_t^* via affecting $(\xi_t)_{t\geq 0}$ which feeds in to $dp_t^* = \beta[d\xi_t - (a^*(p_t^*) + p_t^*)dt]$

- But p_t used to forecast future ξ_s, and hence future values p_s^*, $s \geq t$
  - $(p_t)_{t\geq 0}$ is exogenous

$$\Rightarrow q(p) = \frac{\partial V}{\partial p^*}(p^*, p)\Big|_{p^*=p} \tag{1}$$

- When an equilibrium exists, no belief asymmetry is created: find expression for $q(\cdot)$ when creating such belief asymmetry is sub-optimal

An Expression for $q(\cdot)$

- Proposition 1 in the paper (for $\kappa = 0$ case): If $a^*(\cdot)$ is a Markov equilibrium, $g'(a^*(p)) = \beta q(p)$, where

$$q(p) = \mathbb{E}\left[\int_0^\infty e^{-rt}\left(\chi'(p_t) - g'(a^*(p_t))\left[1 + \frac{da^*}{dp^*}(p_t)\right]dt\right)\middle| p_0 = p\right]$$

  and $dp_t = \beta dZ_t$

- In Holmström (1999), $\chi' = 1$ and a^* is constant, so $da^*/dp = 0$:

$$\Rightarrow q(p) = \mathbb{E}\left[\int_0^\infty e^{-rt}\left[1 - g'(a^*)dt\right]\right] = \frac{1}{r} - \underbrace{\frac{g'(a^*)}{r}}_{=\beta q(p)/r}$$

$$\Rightarrow q(p) = \frac{1}{\beta + r} \qquad (2)$$

Ratcheting and Private Information

$$\Rightarrow q(p) = \underbrace{\frac{1}{r}}_{U'(p)} - \underbrace{\frac{g'(a^*)}{r}}_{\text{correction}=\beta q(p)/r} \qquad \Rightarrow q(p) = \frac{1}{\beta + r}$$

- Correction term is a private-belief driven ratchet effect
  - In eqbm, $p_t = p_t^*$ and $dp_t^* = \beta \underbrace{[d\xi_t - (a^* + p_t^*)dt]}_{\text{unpredictable}} \Rightarrow \mathbb{E}_t[p_s^*] = p_t^*,\ s \geq t$

    $\Rightarrow$ mg. increase in p_t^* yields extra cont. value $\frac{1}{r} = U'(p_t^*)$
  - But if extra effort is used to create abnormally high output, $p_t < p_t^* \Rightarrow$ worker's perspective: market has ratcheted up its expectations of output
  - Replicating extra gain $1/r$ requires more than one-time effort increase: otherwise impact on p^* is less than fully persistent $\Rightarrow$ weaker incentives

Summary

- Pvte info. matters via the channel of good performance potentially leading to a tougher incentive scheme; channel operates via a ratcheting force
  - Strategically affecting beliefs entails a dynamic cost

- The ratchet effect is "hidden" in Holmström because its strength is *uniform* (i.e., independent of (p, p^*)) and of size β

$$\frac{1}{r} \longrightarrow \frac{1}{r+\beta}$$

  - No pvte. info. in eqbm., yet the possibility of its appearance matters

- More generally

$$q(p) = \mathbb{E}\Big[\int_0^\infty e^{-rt}\Big[\chi'(p_t) - \underbrace{g'(a^*(p_t))\Big[1 + \frac{da^*}{dp^*}(p_t)\Big]}_{=q(p_t)\beta[1+da^*/dp^*]:\ \text{ratcheting cost}}\Big]dt\Big]$$

how the market updates its conjecture a^* matters for incentives $\Rightarrow$ strength of ratcheting is endogenous

Rest of the Paper

1. System of ODEs for (q, U) as a necessary condition for Markov equilibria
   - Endogenous ratcheting $\Rightarrow$ q-ODE $\neq$ "Euler equations"

2. Verification theorem: conditions under which a solution to the system in fact constitutes an equilibrium
   - Requires estimating how the "worker" optimally uses his private information + hidden actions
   - Holmström: changes in p do not affect effort ($\beta/(r+\beta)$ still optimal after deviations)

3. Existence of Markov equilibria

4. Applications of ratchet effects in settings with nonlinear payoffs

▸ Ratcheting and Monetary Policy

Optimal Incentive Schemes

Sannikov (2008): Model

- Output process $(\xi_t)_{t\geq 0}$

$$d\xi_t = a_t dt + \sigma dZ_t, \quad \text{where}$$

- As before, ξ_t is total output up to t (so $d\xi_t$ output created over $[t, t+dt)$) and $(Z_t)_{t\geq 0}$ a Brownian motion ($Z_t \sim \mathcal{N}(0, \sigma^2 t)$)

- Agent's flow utility over $[t, t+dt)$ is given by $[u(w_t) - g(a_t)]dt$
  - $u : \mathbb{R}_+ \to \mathbb{R}_+$, $u(0) = 0$, $u' > 0$, $u'' < 0$ and $g : A \to \mathbb{R}_+$ continuous, increasing and convex with $g(0) = 0$ (Technical: $u'(c) \to 0$ as $c \to \infty$ and there is γ_0 s.t. $g(a) \geq \gamma_0 a$ for all $a \in A$)
  - Wage w_t can depend on $\xi^t := (\xi_s : 0 \leq s < t)$ in any way (up to mathematical formalities)

Sannikov (2008): The Principal's Problem

- Expected profits

$$\mathbb{E}^a\left[\int_0^\infty e^{-rs}(d\xi_s - w_s ds)\right] = \mathbb{E}^a\left[\int_0^\infty e^{-rs}(a_s - w_s)ds\right]$$

- Choose $w := (w_t)_{t\geq 0}$ and recommendation $a := (a_t)_{t\geq 0}$ functions of $(\xi^t)_{t\geq 0}$ that solve

$$\begin{aligned}
\max_{(w,a)} \quad & \mathbb{E}^a\left[r\int_0^\infty e^{-rs}(a_s - w_s)ds\right] \\
s.t. \quad & (PC):\ \mathbb{E}^a\left[r\int_0^\infty e^{-rs}[u(w_s) - g(a_s)]ds\right] \geq W_0 \\
& (IC):\ a \in \arg\max_{\tilde{a}}\ \mathbb{E}^{\tilde{a}}\left[r\int_0^\infty e^{-rs}[u(w_s) - g(\tilde{a}_s)]ds\right]
\end{aligned}$$

From Market-Based Incentives

1. $w_t = p_t^* + a_t^*$ or $\chi_t := \chi(p_t^*, a^*(p_t^*))$ are functions of ξ^t via the public belief

$$p_t^* = ... + \beta \int_0^t e^{-\beta s}[d\xi_s - a^*(p_s^*)]ds \longrightarrow p_t^* = p_t^*[\xi^t]$$

   In Holmström (1999) dependence is linear; In Cisternas (2017) need not

2. Given w_t or χ_t, effort characterized by FOC

$$g'(a_t^*) = q(p_t) := \frac{\partial V}{\partial p^*}(p_t, p_t)$$

   i.e., by the sensitivity of the payoff w.r.t. the controlled state (p^*)

   - (that in $q(p_t)$ there is pvte. info. is an orthogonal topic)

Given any compensation scheme, what action would the worker choose?

We need a measure of marginal utility for non-Markov settings

Continuation Value and Sensitivity to Output

- Sannikov uses the **Martingale Representation Thm** (every martingale w.r.t. the Brownian information is an integral against a BM) to obtain such measure of sensitivity/marginal utility
  - MRT also appears in Holmström and Milgrom (1987)
- As before, define the continuation value

$$V_t = \mathbb{E}^a \left[r \int_t^\infty e^{-r(s-t)} [u(w_s) - g(a_s)] ds \middle| (X_s : 0 \leq s < t) \right]$$

- Sannikov shows that given $(w_t, a_t)_{t \geq 0}$ there is $(Y_t)_{t \geq 0}$ s.t.

$$dV_t = \underbrace{r(V_t - [u(w_t) - g(a_t)])dt}_{\text{promise keeping}} + r \cdot \underbrace{Y_t}_{\text{risk exposure}} \cdot \underbrace{[d\xi_t - a_t dt]}_{\text{agent controls } \xi}$$

$$\Rightarrow \text{“}\frac{dV_t}{d\xi_t}\text{”} = rY_t$$

(r is just a normalization)

Incentive Compatibility

- Sannikov (Prop 2): given $(w_t)_{t\geq 0}$, $(a_t)_{t\geq 0}$ is optimal for the worker iff

$$Y_t a_t - g(a_t) \geq Y_t a - g(a) \text{ for all } a \in A \ (*)$$

  where $(Y_t)_{t\geq 0}$ is the sensitivity of $(W_t)_{t\geq 0}$ under (w_t, a_t)

- Put differently, $(a_t)_{t\geq 0}$ interior optimal given $(w_t)_{t\geq 0}$ iff

$$g'(a_t) = Y_t = \frac{dV_t}{d\xi_t}$$

- **One-shot deviation principle**: $\tilde{a}_t$ over $[t, t+dt)$ and a_s, $s > t$
  - Expected gain $rY_t(\tilde{a}_t - a_t)dt$; extra cost $= r[g(\tilde{a}_t) - g(a_t)]$
  - Agent won't deviate if: $Y_t a_t - g(a_t) \geq Y_t \tilde{a}_t - g(\tilde{a}_t)$

  $\Rightarrow$ (*) is an incentive constraint: to implement a_t, principal offers contingent continuation utility "titles" using $Y_t = g'(a_t)$. $(w_t)_{t\geq 0}$ cannot be chosen arbitrarily: linked to Y_t via $(V_t)_{t\geq 0}$

HJB Equation

- Risk is costly: $\gamma(a) := \min\{y \in [0,\infty) : a \in \arg\max_{a' \in A}\ ya' - g(a')\}$

$$dV_t = r(V_t - [u(w_t) - g(a_t)])dt + r\gamma(a) \cdot \underbrace{[d\xi_t - a_t dt]}_{=\sigma dZ_t \text{ as agent follows } a_t}$$

- Value to the principal $F(\cdot)$ then satisfies **ODE**

$$rF(V) = \sup_{a,w}\{a - w + r(V - u(w) + g(a))F'(V) + \frac{1}{2}r^2\gamma(a)^2\sigma^2 F''(V)\}$$

- Agent can always guarantee $V \geq 0$. Why? Boundary conditions
  - $F(0) = 0$: $V = 0 \Rightarrow w_t \equiv 0 \Rightarrow a_t \equiv 0 \Rightarrow F(0) = 0$
  - Ppal can always **retire** the agent: offer constant payment $u^{-1}(V)$ and no production occurs. This costs the principal $F_0(V) = -u^{-1}(V)$

- Look for F over the largest $\mathcal{V} := [0, V_{gp}]$ s.t. $F \geq F_0$ over $\mathcal{V}$ with

$$F(0) = 0,\ \underbrace{F(V_{gp}) = F_0(V_{gp})}_{\text{value matching}} \text{ and } \underbrace{F'(V_{gp}) = F_0'(V_{gp})}_{\text{smooth pasting}}$$

Sannikov's Verification Theorem

Theorem (**Sannikov**, 2008)

1. $\exists!$ *solution* $F \geq F_0$ *to the HJB equation satisfying the boundary conditions for some* $V_{gp} \geq 0$. *Such* F *is concave and it corresponds to the principal's value function.*

2. *For* $V^o \in (0, V_{gp})$, *and as long as* $V_t \in (0, V_{gp})$, $w_t = w(V_t)$ *and* $a_t = a(V_t)$, *with* $a(\cdot)$ *and* $w(\cdot)$ *the maximizers in the HJB, and*

$$dV_t = r(V_t - u(w_t) + g(a_t))dt + r\gamma(a_t)[d\xi_t - a_t dt], \ t > 0, \ V_0 = V^o$$

with $\gamma(a_t) > 0$, *until the time* τ *at which* V_τ *hits* 0 *or* V_{gp}. *In that case, agent gets constant consumption equal to* $-F_0(V_\tau)$ *thereafter.*

Numerical Examples

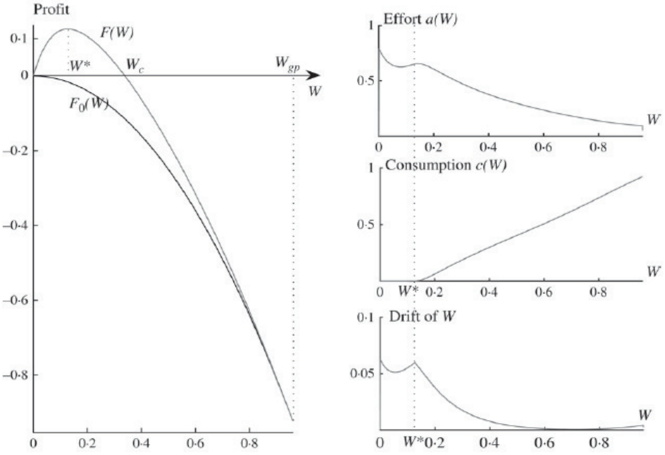

FIGURE 1

Function F for $u(c) = \sqrt{c}$, $h(a) = 0{\cdot}5a^2 + 0.4a$, $r = 0{\cdot}1$ and $\sigma = 1$. Point W^* is the maximum of F

Reference: Sannikov (2008) (variables $(c, h, W) \leftrightarrow (w, g, V)$)

Ratcheting in Monetary Policy

Macroeconomics: Monetary Policy

- Central bank and an economy; a_t: **money growth**

$$\begin{aligned} d\theta_t &= -\kappa\theta_t dt + \sigma_\theta dZ_t^\theta && \textbf{inflation trend - hidden} \\ d\xi_t &= (a_t + \theta_t)dt + \sigma_\xi dZ_t^\xi && \textbf{log price index - public} \end{aligned}$$

- *Phillips curve*: log employment (n) evolves according to

$$dn_t = -\kappa_n n_t dt + \nu \underbrace{[d\xi_t - (a_t^* + p_t^*)dt]}_{\text{unanticipated inflation; } dS_t}$$

- Changes in p_t^* also driven by $dS_t \Rightarrow n_t = p_t^*$ some parameters

- Central bank chooses $(a_t)_{t\geq 0}$ to maximize

$$\mathbb{E}^a\left[\int_0^\infty e^{-rt}\left(-\frac{n_t^2}{2} - \frac{a_t^2}{2}\right)dt\right]$$

Equilibrium inflation? Flow payoff **nonlinear** in $n_t = p_t^*$

Monetary Policy

- Choose $\kappa_n = \kappa$, $\nu = \beta$ and $n_0 = p_0$, so

$$dn_t = -\kappa n_t dt + \beta[d\xi_t - (a_t^* + p_t^*)dt], \; n_0 = p_0$$

$\Rightarrow n_t = p_t^*$

- Suppose $(\theta_t)_{t \geq 0}$ is observable or absent. Then

$$dn_t = [-\kappa n_t + \beta(a_t - a_t^*)]dt + \sigma dZ_t^\xi$$

$\Rightarrow$ environment becomes one of imperfect monitoring only

Proposition (Observable case)

In any linear equilibrium, $a^{*,o}(n) = \beta \alpha^o n$*, where* $\alpha^o < 0$*.*

Monetary Policy

- Hidden case: Phillips curve now becomes

$$dn_t = [-\kappa n_t + \beta(a_t - a_t^*) - \underbrace{\beta(p_t^* - p_t)}_{\text{ratcheting}}]dt + \sigma dZ_t$$

But $a_t > a_t^* \Rightarrow p_t^* > p_t$, i.e. ratcheting puts **extra downward** pressure on employment

Proposition (Hidden case)

There exists a linear eqbm. $a^{*,h}(n) = \beta\alpha^h n$, $\alpha^h < 0$, *s.t.* $|\alpha^h| < |\alpha^o|$.

Monetary Policy

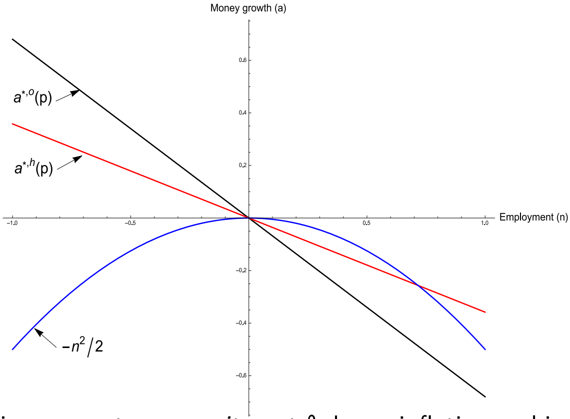

- Ratcheting generates *commitment* & lower inflationary bias for $n < 0$

$$\underbrace{\left[r + 2\kappa + \beta^2\alpha^o\right]\alpha^o = -1}_{\text{obs. case}} \quad \text{and} \quad \underbrace{\left[r + 2\kappa + \beta + \beta^2\alpha^h\right]\alpha^h = -1}_{\text{hidden case}},$$

Back to Rest of the Paper