Dynamic Contracts:

A Continuous-Time Approach

Yuliy Sannikov

Stanford University

Plan

- Background: principal-agent models in economics
  - Examples: what questions are economists interested in?
- Continuous-time approach to dynamic agency
  - Basic model – double optimization
  - Agent’s incentives
  - Principal’s control problem
  - Characteristics of solution
  - Going further: persistent private information, hidden savings, etc
- Open questions

Principal-agent models in economics

- Corporate finance
  - Capital structure and financing constraints
  - Executive compensation
- Macroeconomics / public finance
  - Taxation / insurance / incentives
- Personnel economics
  - Labor contracts
- Industrial organizations

Basic model

- Phelan and Townsend (1991): time t = 0, 1, 2, …
- Probability of output given effort: $P(x \mid a)$
- In period t, principal observes $x^t = \{x_0, x_1, \dots x_t\}$, but not $\{a_0, a_1, \dots a_t\}$, pays the agent $c_t(x^t)$
- Problem: design contract $\{c_t(x^t)\}$ to maximize

$$E\left[\sum_{t=0}^{\infty} \delta^t (x_t - c_t)\right]$$

$$\text{s.t. } E\left[\sum_{t=0}^{\infty} \delta^t u(a_t, c_t)\right] = w$$

$$\{a_t\} \text{ maximizes } E\left[\sum_{t=0}^{\infty} \delta^t u(a_t, c_t)\right] \text{ given } c_t(x^t)$$

Optimal taxation: the Mirrlees model

- Taxpayers' abilities ~ $f(\theta)$, θ is wage, choose labor l
- Gov't observes only income θl, but not θ, l separately
- Agents get utility $u(c(\theta), l)$, $c(\theta) = \theta l - T(\theta l)$
- Choose tax policy $T()$ to maximize

$$\text{Social welfare function } \int G\big(u(c(\theta), l(\theta))\big) f(\theta) d\theta$$

$$\text{s.t. (budget)} \quad \int T\big(\theta l(\theta)\big) f(\theta) d\theta = R$$

$$(IC) \quad l(\theta) = \arg\max u(\theta l - T(\theta l), l)$$

- Dynamic models (Werning, Farhi, Golosov, Tsyvinski) – stochastic process for ability θ, investment in human capital, etc.

Dynamic price discrimination

- Battaglini (AER 2005)
- Buyer's per-period utility: $\theta_t q_t - p_t$
- Seller's cost: $q_t^2/2$
- Buyer's type θ_t: Markov process (private info)
- In period t, seller offers any q at price $p(q)$ (schedule depends on entire past history)
- Objective: maximize profit

$$E\left[\sum_{t=0}^{\infty} \delta^t (p_t - q_t^2 / 2)\right]$$

s.t. (1) buyer's expected utility ≥ 0 and

(2) sequence q_t maximizes buyer's utility given pricing policy and type sequence θ_t

Basic Theory

- **Discrete Time:** Radner (*EMA* 1985), Rogerson (*EMA* 1985), Spear and Srivastava (*ReStud* 1987), Fudenberg, Holmstrom and Milgrom (*JET* 1990), Phelan and Townsend (*ReStud* 1991)
  - Patient agent $\rightarrow$ efficiency is attainable
  - Optimal contract is recursive (agent's continuation value is the state variable that controls incentives)

- **Why continuous time?** Discrete-time models are messy – full of details that distract from big picture
  - E.g. in Phelan and Townsend (1991), optimal contract involves randomizations – before the agent puts in effort, principal randomizes over which effort he wants to incentivize
  - Agent's payoff follows a random walk with many step sizes – as many as output levels – distracts from long-term distribution properties

Applications (using continuous-time approach)

- DeMarzo and Sannikov (JF 2006): corporate finance application, unobservable cash flows, capital structure question
- Biais, Mariotti, Plantin and Rochet (ReStud 2007), He (JFE 2007), Biais, Mariotti, Rochet and Villeneuve (EMA 2010), Hoffmann and Pfeil (RFS 2010), Piskorski and Tchistyi (RFS 2010, 2011), DeMarzo, Fishman, He and Wang (JF 2012), DeMarzo and Sannikov (ReStud 2017), He, Wei, Yu and Gao (RFS 2017)

- Optimal taxation / insurance: Farhi and Werning (ReStud 2013), Williams (EMA 2011)

- Repeated Games: Sannikov (*EMA* 2006)
- Games with reputation: Faingold and Sannikov (*EMA* 2011)

Basic Model (continuous time)

- Time $t \in [0, \infty)$
- Risk-neutral principal and risk-averse agent, common discount rate *r*
- Agent puts in effort $A = \{A_t \in \mathcal{A}, 0 \leq t < \infty\}$
- Principal does not see effort, but observes output

$$dX_t = A_t \, dt + \sigma dZ_t$$

where Z is a Brownian motion

- Cost of effort $h : \mathcal{A} \to \Re$ continuous, increasing, convex with $h(0) = 0$.
- Utility of consumption $u : [0, \infty) \to [0, \infty)$ continuous, increasing and concave with $u(0) = 0$ and $u'(c) \to 0 \text{ as } c \to \infty$

Based on Sannikov (2008) "A Continuous-Time Version of the Principal-Agent Problem," *ReStud*

Problem Formulation

Find a contract $\{C_t, 0 \leq t < \infty\}$ and effort recommendation $A = \{A_t \in \mathcal{A}, 0 \leq t < \infty\}$ that maximizes the principal's profit

$$E^A\left[r \int_0^\infty e^{-rt} (dX_t - C_t)\, dt \right]$$

subject to

$$W_0 = E^A\left[r \int_0^\infty e^{-rt} (u(C_t) - h(A_t))\, dt \right]$$

and $$W_0 \geq E^{\hat{A}}\left[r \int_0^\infty e^{-rt} (u(C_t) - h(\hat{A}_t))\, dt \right]$$

Nature of the problem

- **Two embedded** dynamic optimization problems
  - Principal is offering a dynamic contract recognizing that the agent will be optimizing dynamically

- **To solve:**
  - introduce new process – agent's continuation payoff – to characterize the agent's incentives
  - this reduces the principal's problem to optimal stochastic control

4 steps to derive the optimal contract

Principal's problem: optimal stochastic control

1. Define the agent's continuation value $(W_t)_{t \geq 0}$ for any

$$\{C_t, 0 \leq t < \infty\} \qquad A = \{A_t \in \mathcal{A}, 0 \leq t < \infty\}$$

2. Using the Martingale Representation Theorem, write an equation that describes the evolution of W_t
3. Necessary and sufficient conditions for the agent's effort to be optimal (sensitivity of W_t to X_t)
4. Solve the principal's problem using the HJB equation

The Agent's Continuation Value W_t

- **Step 1:** Given consumption $\{C_s, 0 \leq s < \infty\}$ and effort $\{A_s, 0 \leq s < \infty\}$, define the agent's continuation value

$$W_t = E_t\left[r \int_t^{\infty} e^{-r(s-t)} (u(C_s) - h(A_s)) ds \right]$$

The Agent's Continuation Value W_t

- **Step 2:**

Proposition 1. In any contract with finite payoff to the agent, W_t is the agent's continuation value if and only if

$$dW_t = r(W_t - u(C_t) + h(A_t))dt + rY_t(dX_t - A_t dt).$$

for some process $\{Y_t\}$ and $E[e^{-rt} W_t] \to 0$.

Sketch of proof. Existence of Y_t follows from the representation of

$$r\int_0^t e^{-rs}(u(C_s) - h(A_s))ds + e^{-rt}W_t = E_t\left[r\int_0^\infty e^{-rs}(u(C_s) - h(A_s))ds\right]$$

using Martingale Representation Theorem.

Incentives

The agent maximizes

depends on $\hat{A}_t$

$$E[r(u(C_t) - h(\hat{A}_t))dt + dW_t]$$

$$dW_t = r(W_t - u(C_t) + h(A_t))dt + rY_t(\boxed{dX_t} - A_t dt)$$

Agent will maximize $\quad Y_t \hat{A}_t - h(\hat{A}_t)$

Step 3:

Proposition 2. A contract is incentive-compatible if and only if

$$\forall t \geq 0,\ A_t \text{ maximizes } Y_t a - h(a) \qquad \text{(IC)}$$

If (IC) holds, we say Y_t enforces A_t

Sketch of Proof

Consider $$\hat{V}_t \equiv r\int_0^t e^{-rs}(u(C_s) - h(\hat{A}_s))\,ds + e^{-rt}W_t$$

If (IC) holds, then for any deviation, $\hat{V}_t$ is an $\hat{A}$-supermartingale

Hence, $$W_0 = \hat{V}_0 \geq E^{\hat{A}}\left[r\int_0^t e^{-rs}(u(C_s) - h(\hat{A}_s))\,ds + e^{-rt}W_t\right],$$

and taking t to infinity, we find that A is not worse than $\hat{A}$

If (IC) fails, we can find a deviation such that $\hat{V}_t$ is an $\hat{A}$-submartingale. Hence,

$$W_0 = \hat{V}_0 < E^{\hat{A}}\left[r\int_0^t e^{-rs}(u(C_s) - h(\hat{A}_s))\,ds + e^{-rt}W_t\right]$$

The Optimal Control Problem

Propositions 1 and 2 imply:

Theorem: There is a one-to-one correspondence between

- Contracts $\{C_t, t \geq 0\}$ with strategies $\{A_t, t \geq 0\}$ that satisfy the incentive constraints, with finite value to the agent and

- Controlled processes

$$dW_t = r(W_t - u(C_t) + h(A_t))dt + rY_t\sigma\, dZ_t$$

that satisfy the transversality condition $\lim_{t\to\infty} E[e^{-rt}W_t] = 0$, with controls $\{C_t, A_t, Y_t\}$ such that Y_t enforces A_t

The Optimal Control Problem

$$dW_t = r(W_t - u(C_t) + h(A_t))dt + rY_t\sigma\, dZ_t$$

- The principal
  - controls W_t with C_t, A_t and Y_t (which enforces A_t)
  - must honor promises, i.e. $E[e^{-rt}W_t] \to 0$
  - gets expected flow of profit of $A_t - C_t$

Denote by $F(W_0)$ the maximal total profit that the principal can attain in this way

HJB equation

Controlled process:

$$dW_t = r(W_t - u(C_t) + h(A_t))dt + rY_t\sigma\, dZ_t$$

HJB equation:

$$\boxed{rF(W) = \max_{\substack{c,a,Y \text{ s.t.} \\ Y \text{ enforces } a}} r(a-c) + r(W - u(c) + h(a))F'(W) + \frac{rY^2\sigma^2}{2}F''(W)}$$

Denote by y(a) minimal (in absolute value) Y that enforces a

Profit function F vs. first best

Let $F_0(u(c)) = -c$ be "retirement" profit. Solve

$$F''(W) = \min_{a>0,c} \frac{F(W) + c - a - (W - u(c) + h(a))F'(W)}{r\gamma(a)^2\sigma^2/2}$$

with $F(0) = 0$, and largest $F'(0)$ such that $F(W^{gp}) = F_0(W^{gp})$ for some $W^{gp} \geq 0$

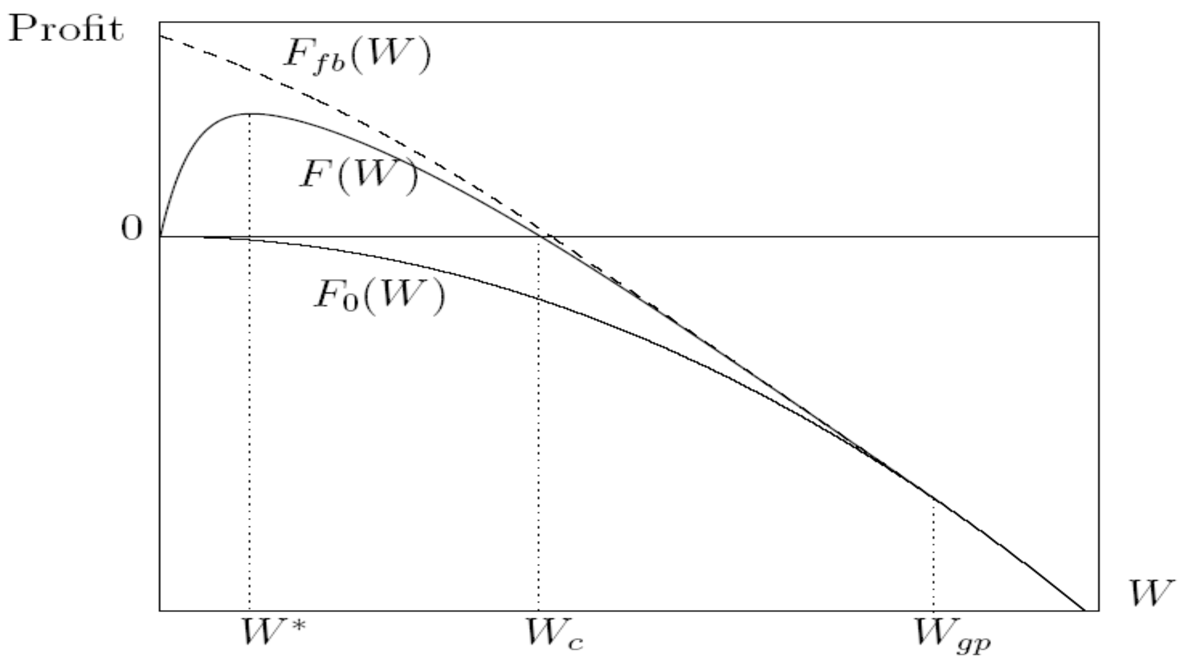

The Optimal Contract

$F(W_0)$ is the principal's profit in the optimal contract for $W_0 \in [0, W^{gp}]$. The agent's value in the optimal contract follows

$$dW_t = r\left(W_t - u(c(W_t)) + h(a(W_t))\right)\, dt + rY(W_t)\, \underbrace{(dX_t - a(W_t)dt)}_{\sigma dZ_t},$$

until the retirement time τ when W_t hits 0 or W^{gp}. For $t < \tau$ $C_t = c(W_t)$ and $A_t = a(W_t)$ are the maximizers in the ODE for $F(W)$. After time τ, the agent receives constant consumption $C_t = -F(W_\tau)$ and puts effort 0.

An Example

$$u(c) = \sqrt{c},\ h(a) = 0.5a^2 + 0.4a,\ r = 0.1 \text{ and } \sigma = 1.$$

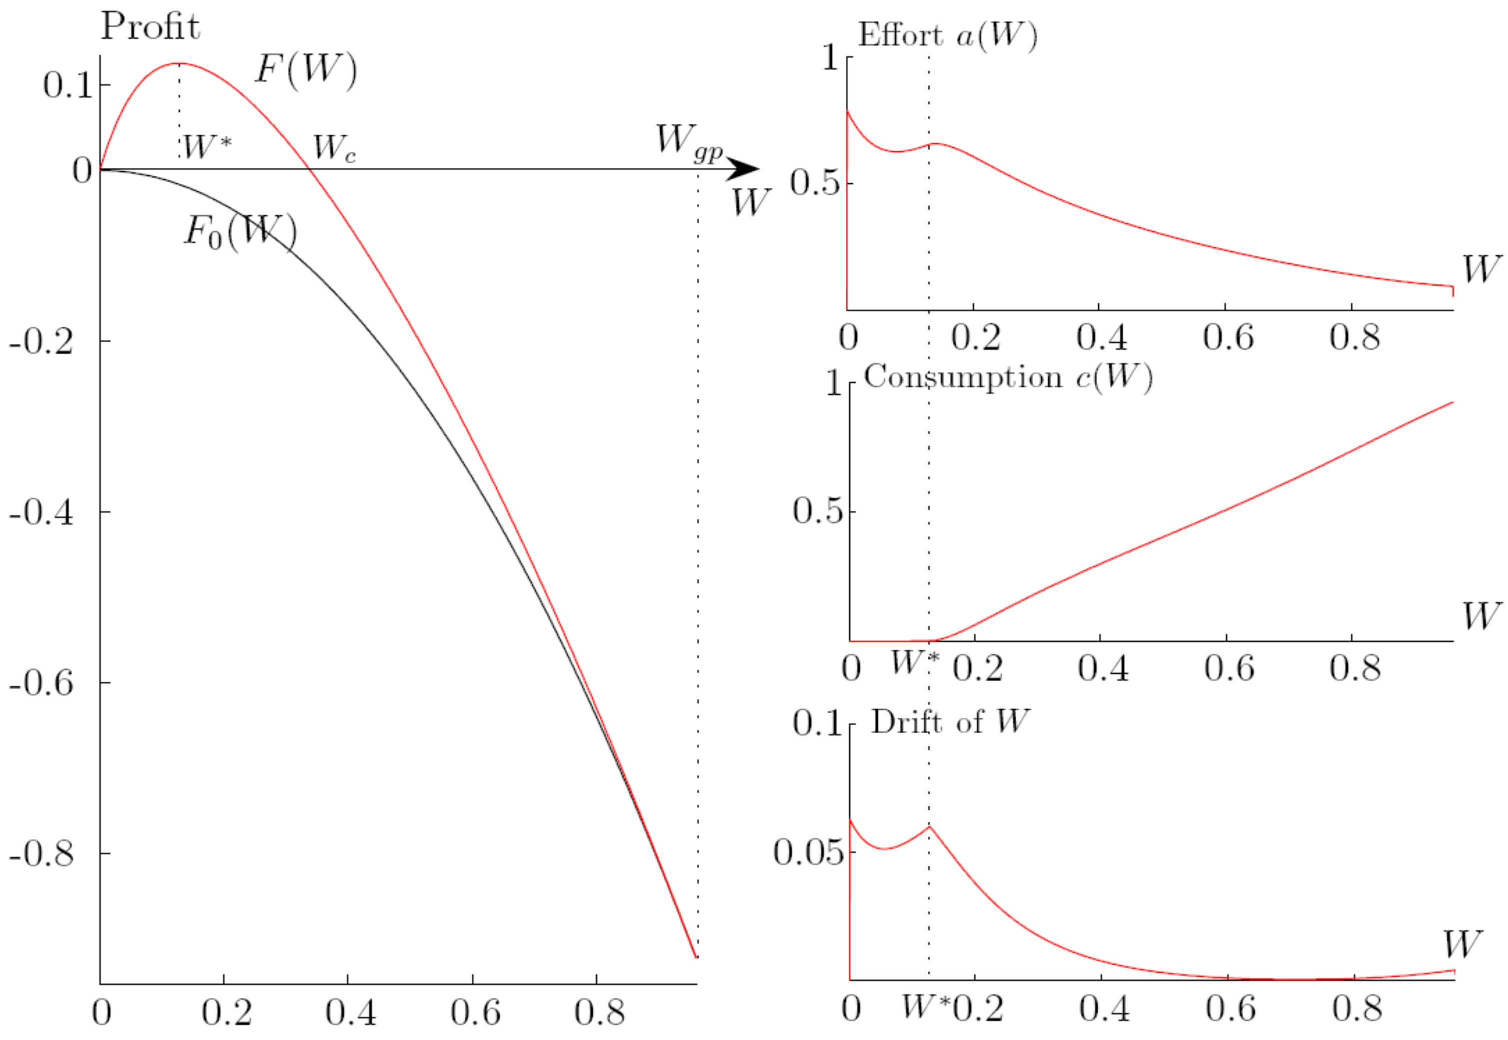

Properties of solution

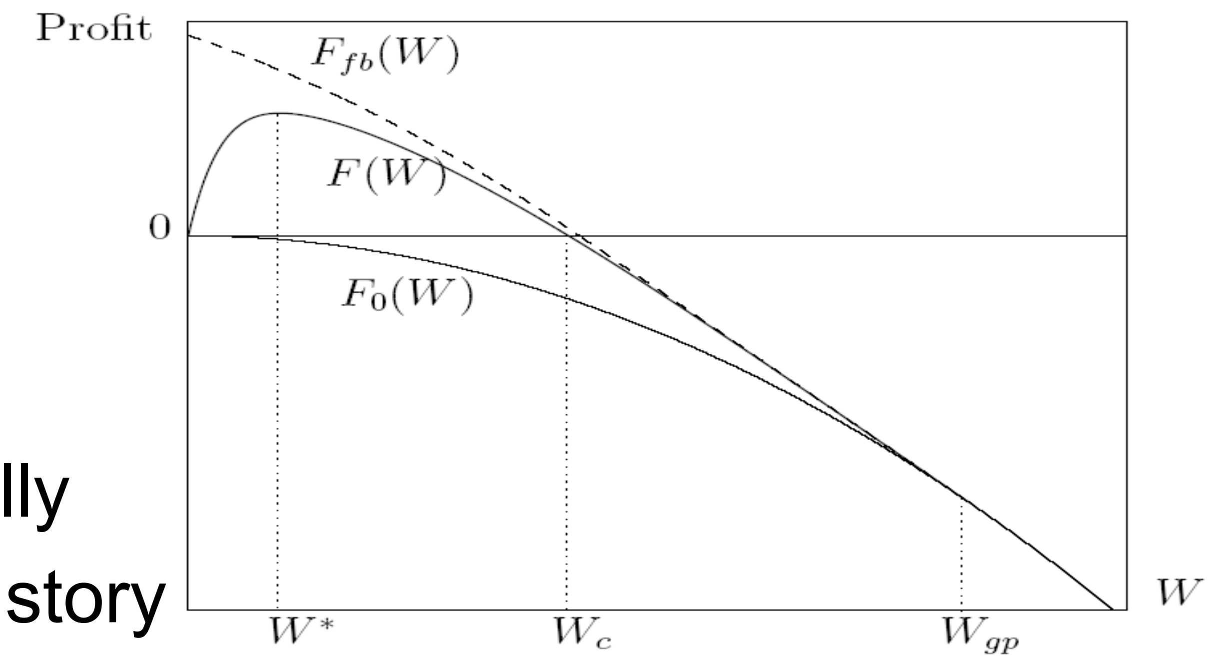

- Cont. value W_t fully summarizes past history
  - bad performance history can be undone by good performance (unless the agent is retired at 0)
- It's optimal to eventually let agent retire (at 0/W^{gp})
- Contract is not renegotiation-proof
  - *F(W)* has an increasing portion: principal punishes the agent at a cost
- Payoff W^* that maximizes *F(W)* is positive
  - principal needs room to reward and punish the agent

Where can we go from here?

- Easy: change boundary conditions
  - The agent's outside option
  - The cost of replacing the agent
  - Promotion opportunities
- Harder: need to add more state variables to summarize the agent's incentives
  - Agent's effort affects output now and in the future
  - Agent privately observes persistent shocks to output
  - Hidden savings

Example: asset management + hidden savings

- Agent manages capital k_t, obtains return per dollar

unobservable, ≥ 0, "stealing"

observable → $$dR_t = (\alpha + r - a_t)\,dt + \sigma\, dZ_t$$

- Agent's utility

$$W^{a,\hat{c}} = E^a\left[\int_0^\infty e^{-rt}\frac{\hat{C}_t^{1-\gamma}}{1-\gamma}\,dt\right]$$

- $a_t \geq 0$ is "stealing". Hidden savings, $h_t \geq 0$

$$dh_t = \left(rh_t + C_t - \hat{C}_t + \phi k_t a_t\right)dt, \qquad \phi \leq 1$$

- Principal specifies C_t and k_t given history of returns

Based on joint work with Sebastian Di Tella (Stanford GSB)

Incentives to not steal

Optimal contract involves no stealing (stealing is inefficient) and no savings (without loss of generality)

Sensitivity of agent's expected payoff to returns

$$dW_t = r(W_t - \tfrac{C_t^{1-\gamma}}{1-\gamma})dt + r\Delta_t(dR_t - (r+\alpha)dt)$$

Incentive constraint:

$$\Delta_t \geq C_t^{-\gamma}\phi k_t$$

Remark: C_t affects incentives but depends on the precautionary motive

Incentives to not save

- Agent has no incentives to save when marginal utility

$$C_t^{-\gamma}$$

is a supermartingale (constant in expectation – or decreasing)

- Conditions on W_t and C_t rule out profitable deviations with stealing or saving (first-order approach)
- What about double deviations? Steal, save, consume later. We have to keep our fingers crossed…

Properties of solution

- Control problem, 2 state variables: *W* and *C*
- Ratio *C/X*, where

$$X_t = ((1-\gamma)W_t)^{\frac{1}{1-\gamma}}$$

is "utility in dollars" reflects the agent's precautionary motive.

  - C/X goes down when precautionary motive increases
- Given W_0, the principal sets C_0/X_0 to maximize profit
- But $C_t/X_t > C_0/X_0$ for $t > 0$: principal distorts the contract to safety to reduce precautionary motive and improve incentives ex-ante
- C_t/X_t goes up after <u>bad</u> outcomes
  - contract gets safer after bad outcomes – reduces precautionary motive

Verifying agent's incentives

- We solved the principal's problem subject to less than full set of incentive constraints by the agents (deviations with only "stealing" and savings)
- Fingers crossed – we are done only if agent does not want to deviate in any other way
- We can verify if we can find an upper bound on the agent's *deviation* value function U(W, C, h) such that

$$U(W, C, h) = W$$

Verifying agent's incentives

- We solved the principal's problem subject to less than full set of incentive constraints by the agents (deviations with only "stealing" and savings)
- Fingers crossed – we are done only if agent does not want to deviate in any other way
- Verification:

$$U(W,C,h) = \left(1 + \frac{hC^{-\gamma}}{(1-\gamma)W}\right)^{1-\gamma} W$$

is an upper bound of C/X goes up (contract gets safer) after bad performance.

- Intuition: if precautionary motive is reduced after bad performance, stealing and saving is unattractive

Some general properties

- Settings with hidden savings, private info about fundamentals, long-term consequences of actions:
- Optimal contract has distortions – principal commits to manage the agent's future "information rents" to improve incentives ex-ante
- No reversibility – bad early performance cannot be undone by future good performance and vice versa

Some general properties

- Harder questions: need to add more state variables to summarize the agent's incentives
  - Agent's effort affects output now and in the future
  - Agent privately observes persistent shocks to output
  - Hidden savings
- Optimal contract has distortions – principal commits to certain ex-post inefficiencies to improve incentives ex-ante
- These distortions are history-dependent – irreversibility – bad early performance cannot be undone by future good performance

Conclusions

- Continuous-time methods offer huge potential to analyze dynamic agency models, which are common in economics
- A lot of progress in the last decade, especially in corporate finance, but also in other fields
- Three potential areas for future fruitful research
  1. Effective numerical methods for solving control problems.
  2. Complexity of problems with persistent private information (like hidden savings) – when does the first-order approach work in general, and what to do when it does not work
  3. Embedding models of dynamic contracts in broader macro settings – determination of interest rates, investment, business cycles, resource constraints