Rules without Commitment: Reputation and Incentives

Alessandro Dovis
U Penn and NBER

Rishabh Kirpalani
Wisconsin

November 2019
LAEF, UC Santa Barbara

This Paper

- Rules often proposed as solution to time inconsistency problem
  - Society can credibly impose rules on policy makers
- In reality substantial *uncertainty* about whether policy makers can resist temptation to deviate
- **How to design rules when there is uncertainty about the ability of policy makers to enforce the rule ex-post?**

Our Approach

- *Rule designer* chooses rules (policy recommendation)
  - EU design fin. regulation, gov't chooses central bank's mandate
- *Policy maker* implement policy
  - Single Resolution Board/ Central banker
- Policy maker can be one of two (hidden) types
  - Commitment type: always follows rule
  - Optimizing type: chooses policy sequentially

  Reputation = probability policy maker is commitment type
- *Private agents* make decisions given
  - Announced rule
  - Expectations about whether rule will be followed

Our Approach

- *Rule designer* chooses rules (policy recommendation)
  - EU design fin. regulation, gov't chooses central bank's mandate
- *Policy maker* implement policy
  - Single Resolution Board/ Central banker
- Policy maker can be one of two (hidden) types
  - Commitment type: always follows rule
  - Optimizing type: chooses policy sequentially

  Reputation = probability policy maker is commitment type
- *Private agents* make decisions given
  - Announced rule
  - Expectations about whether rule will be followed
- **Study optimal rule design problem**
  - **Provide incentives to private agents and policy makers**

More Lenient Rules to Preserve Reputation

- If reputation low: preserve uncertainty over time
  - Most stringent rule that will be followed by optimizing type
  - Rules more lenient than in static setting
  - Uncertainty about policy maker's type beneficial
    - Decreasing returns to reputation

- If reputation high: separation
  - Same as static setting
  - Dynamic losses (uncertainty beneficial) but static benefits

Role for Opaque Rules

- When policy maker's type known:
  Ability to monitor policy makers supports better outcomes
  - Atkeson-Chari-Kehoe (2007) and Piguillem-Schneider (2017)

- When policy maker's type uncertain and reputation high:
  Opaque rules optimal
  - Want rules hard to monitor (hard to detect deviation)
  - Complicated rules contingent on irrelevant contingencies

  Preserve uncertainty w/out static losses associated w/ leniency

Policy Game

Environment

- $t = 0, 1, ..., T$
- Rule designer
- Private agents
- Policy maker
  - Commitment type
  - Optimizing type
- Common prior of commitment type is ρ (reputation)

Timing

- Rule designer announces a rule
  - A rule is a policy recommendation $\pi_r \in [\underline{\pi}, \overline{\pi}]$

- Private agents take their action $x = \phi(\mathbb{E}\pi)$

- Policy maker chooses policy π
  - Commitment type always follows recommendation: $\pi = \pi_r$
  - Optimizing type chooses its policy sequentially : $\pi \in [\underline{\pi}, \overline{\pi}]$

- Social welfare function: $w(x, \pi)$

Timing

- Rule designer announces a rule
  - A rule is a policy recommendation $\pi_r \in [\underline{\pi}, \overline{\pi}]$
- Private agents take their action $x = \phi(\mathbb{E}\pi)$, $\phi' < 0$
- Policy maker chooses policy π
  - Commitment type always follows recommendation: $\pi = \pi_r$
  - Optimizing type chooses its policy sequentially : $\pi \in [\underline{\pi}, \overline{\pi}]$
- Social welfare function: $w(x, \pi)$, $w_x > 0, w_{x\pi} < 0$
- $\pi < \pi' \iff \pi$ more stringent than π'

Time Inconsistency

- Let $(x_{ramsey}, \pi_{ramsey})$ be the Ramsey outcome:

$$(x_{ramsey}, \pi_{ramsey}) = \arg\max_{x,\pi} w(x, \pi) \quad \text{subject to} \quad x = \phi(\pi)$$

- Normalize $\pi_{ramsey} = \underline{\pi}$

- Let $\pi^*(x)$ be best response to x

$$\pi^*(x) = \arg\max_{\pi} w(x, \pi)$$

- Assume that the Ramsey policy is not time-consistent:

$$\pi_{ramsey} = \underline{\pi} < \pi^*(x_{ramsey})$$

Examples

- Barro-Gordon ◂ Barro-Gordon
  - x: wage inflation set by unions
  - π: inflation rate
- Bank-Bailout a la Kareken-Wallace ◂ Bailout
  - x: bankers' effort
  - π: bailout to lenders
- Capital taxation ...

The Benefits of Lenient Rules

Statically Optimal Rule

The optimal rule solves

$$W_0(\rho) = \max_{\pi_c, \pi_o, x} \rho w(x, \pi_c) + (1 - \rho) w(x, \pi_o)$$

subject to

$$x = \phi(\rho \pi_c + (1 - \rho) \pi_o)$$
$$\pi_o = \pi^*(x)$$

Statically Optimal Rule

The optimal rule solves

$$W_0(\rho) = \max_{\pi_c, \pi_o, x} \rho w(x, \pi_c) + (1-\rho) w(x, \pi_o)$$

subject to

$$\begin{aligned} x &= \phi(\rho \pi_c + (1-\rho)\pi_o) \\ \pi_o &= \pi^*(x) \end{aligned}$$

- Solution: $(\pi_0(\rho), x_0(\rho))$
- Assume that $\pi_0(\rho) = \underline{\pi}$ (holds in our examples)
- Value for the optimizing type:

$$V_0(\rho) = w(x_0(\rho), \pi^*(x_0(\rho))).$$

Uncertainty Beneficial

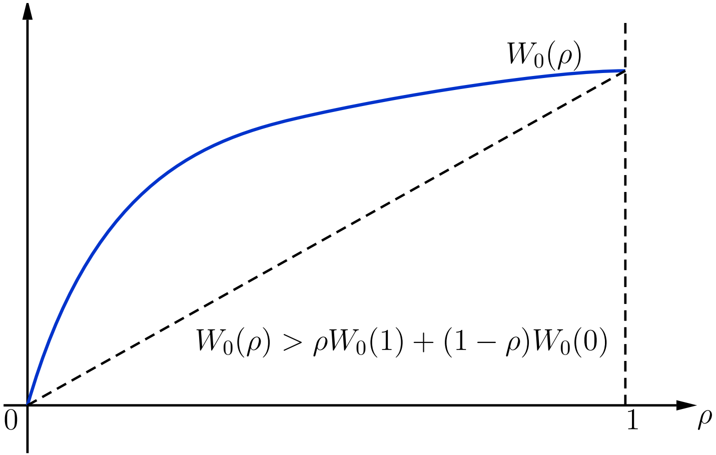

Sufficient condition: $W_0(\rho)$ is concave

When is Uncertainty Beneficial?

Sufficient conditions for $W_0(\rho)$ to be concave

- $w(x, \pi)$ jointly concave
- $w_\pi(x, \pi)$ is convex
- ϕ is concave
- $1 > \pi^*_x(x)\phi'(\pi) > \underline{A}$
- $\pi_c(\rho) = \underline{\pi}$

Satisfied by a large class of environments including both examples

Intuition

- There are decreasing returns to reputation
  - Reputation is more valuable when there is less of it
  - Increase in reputation increases x more when reputation in low

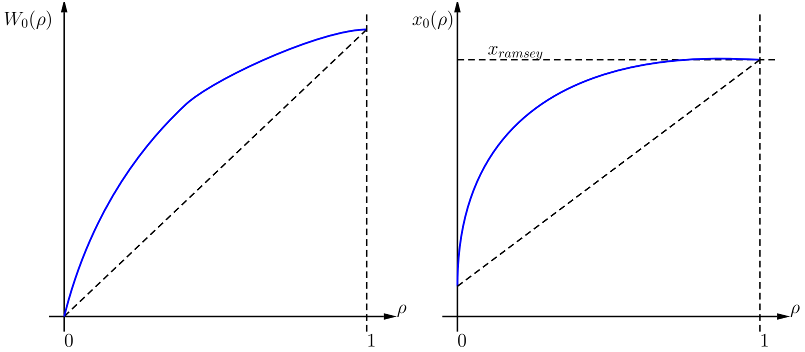

Intuition

- There are decreasing returns to reputation
  - Reputation is more valuable when there is less of it
  - Increase in reputation increases x more when reputation in low

- Recall $x = \phi\left(\rho\underline{\pi} + (1-\rho)\,\pi^*(x)\right)$
  - An increase in ρ increases x through a direct and indirect channel

- x is concave if
  - ϕ is concave
  - π^* is convex (follows from $w_\pi(x,\pi)$ convex)

- Concavity of w and $x(\rho) \Rightarrow w(x,\underline{\pi})$ and $w(x,\pi^*(x))$ are concave

- Show $W_0(\rho) = \rho w(x,\underline{\pi}) + (1-\rho)w(x,\pi^*(x))$ concave

Dynamically Optimal Rule (Twice Repeated)

The optimal rule solves

$$W(\rho) = \max_{x,\pi_c,\pi_o} \rho\left[w(x,\pi_c) + \beta W_0\left(\rho'_c\right)\right] + (1-\rho)\left[w(x,\pi_o) + \beta W_0\left(\rho'_o\right)\right]$$

subject to

$$x = \phi\left(\rho\pi_c + (1-\rho)\pi_o\right),$$

the incentive compatibility constraint for the optimizing type,

$$w(x,\pi_c) + \beta_o V_0\left(\rho'_o\right) \geqslant w(x,\pi^*(x)) + \beta_o V_0(0),$$

and the law of motion for beliefs,

$$\rho'_c = \begin{cases} 1 & \text{if } \pi_o \neq \pi_c \\ \rho & \text{o/w} \end{cases}, \quad \rho'_o = \begin{cases} 0 & \text{if } \pi_o \neq \pi_c \\ \rho & \text{o/w} \end{cases}.$$

Optimizing Type Never Randomizes

- Suppose optimizing type mixes between rule and best response
- This is welfare dominated by case in which it follows rule for sure
- Two reasons
  - Introduces volatility in posterior without affecting its mean
    - Lowers continuation value since uncertainty beneficial
  - Tightens the optimizing type's incentive constraint
    - Since w is concave in π and V_0 is concave in ρ

Main Result

Suppose β_o is small enough

- Ramsey outcome is not IC for the optimizing type for all ρ

Proposition

Under conditions for which uncertainty is beneficial:

- *For ρ close to 1, it is optimal to separate,* $\pi_c \neq \pi_o = \pi^*(x)$
  - *Same outcome as in static setting,* $\pi_c = \underline{\pi}$
- *For ρ close to 0, it is optimal to pool*
  - *Rule is less stringent than in static setting,* $\pi_c = \pi_o > \underline{\pi}$
  - *Most stringent rule consistent with IC for optimizing type*

Pooling vs. Separation

- If there is separation:
  - First period outcome solves static problem:
    - $\pi_c = \pi_0 = \underline{\pi}$ and $\pi_o = \pi^*(x_0)$
  - Expected continuation value is $\rho W_0(1) + (1-\rho) W_0(0)$

- If there is pooling:
  - $\pi_o = \pi_c = \pi_{ico}(\rho)$ most stringent rule consistent with IC for optimizing type

    $$w(x_{ico}, \pi_{ico}) + \beta_o V_0(\rho) = w(x_{ico}, \pi^*(x_{ico})) + \beta_o V_0(0)$$

    where $x_{ico}(\rho) = \phi(\pi_{ico}(\rho))$
  - Expected continuation value is $W_0(\rho)$

Pooling vs. Separation

Define

- Dynamic benefits of pooling

$$\Delta\Omega(\rho) \equiv W_0(\rho) - [\rho W_0(1) + (1-\rho) W_0(0)]$$

  - Always positive since uncertainty is beneficial

- Static benefits of pooling

$$\Delta\omega(\rho) \equiv w(x_{ico}(\rho), \pi_{ico}(\rho)) - W_0(\rho)$$

  - + Optimizing type follows tougher policy (closer to Ramsey)
  - − Commitment type follows lenient policy (further from Ramsey)

Optimal to pool iff

$$\Delta\omega(\rho) + \beta\Delta\Omega(\rho) \geqslant 0$$

Dynamic Benefits

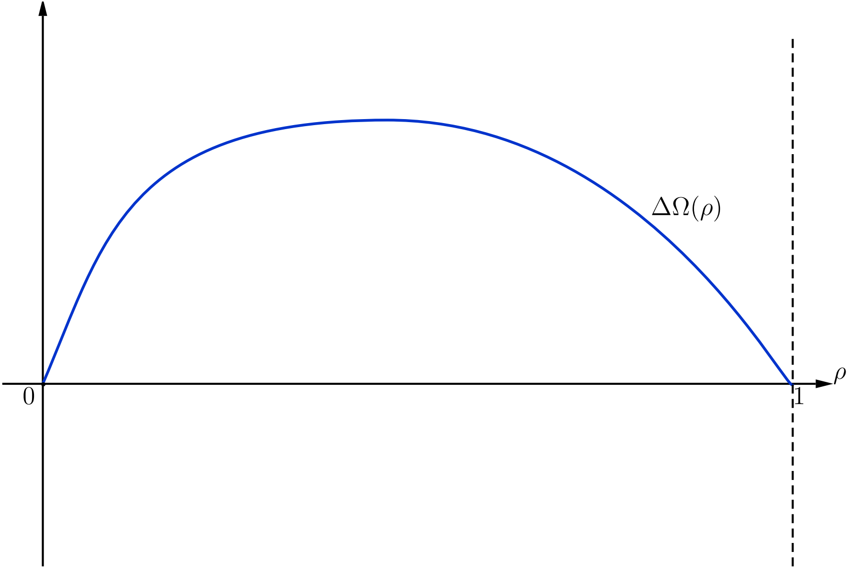

- $\Delta\Omega\,(\rho) \geqslant 0$ with $\Delta\Omega\,(0) = \Delta\Omega\,(1) = 0$

Static Benefits

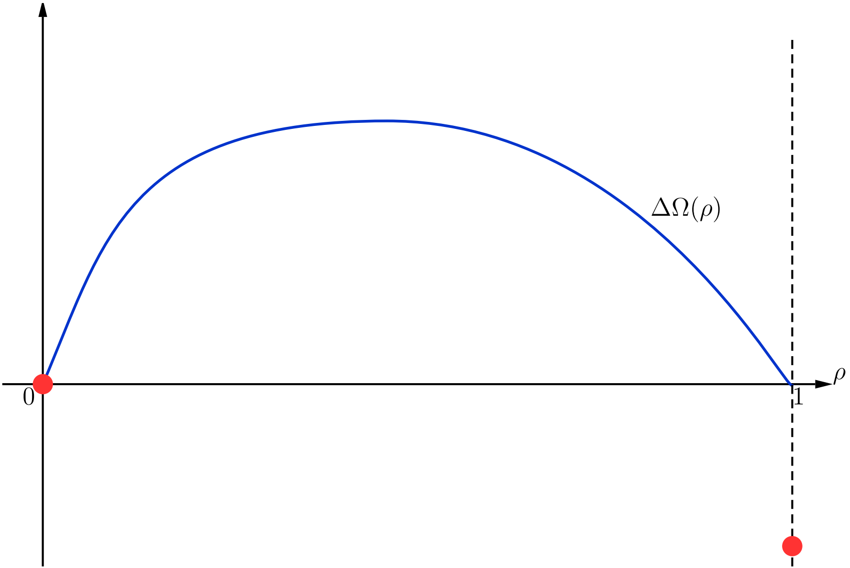

- $\Delta\omega\,(0) = 0$ and $\Delta\omega\,(1) < 0$

Static Benefits

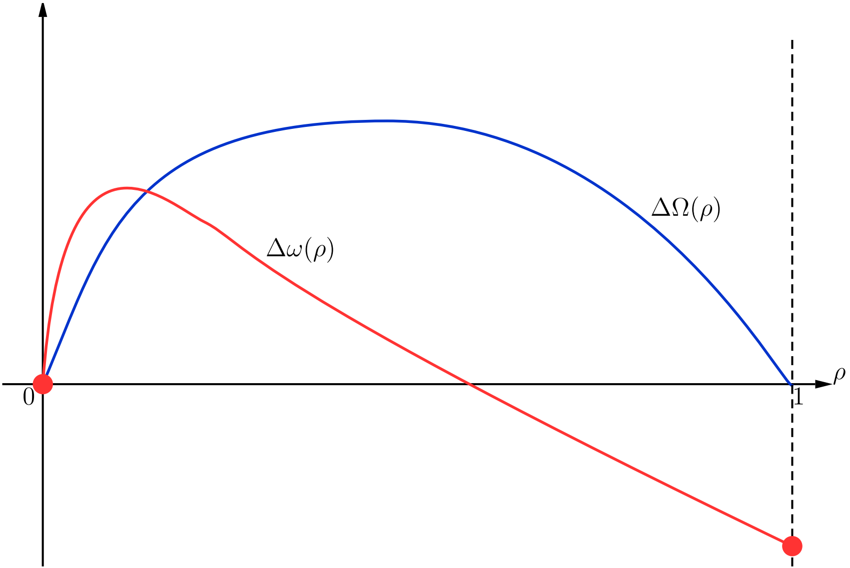

- $\Delta\omega(\rho) > 0$ for ρ close to zero

Low Reputation ⇒ Lenient Rule with Pooling

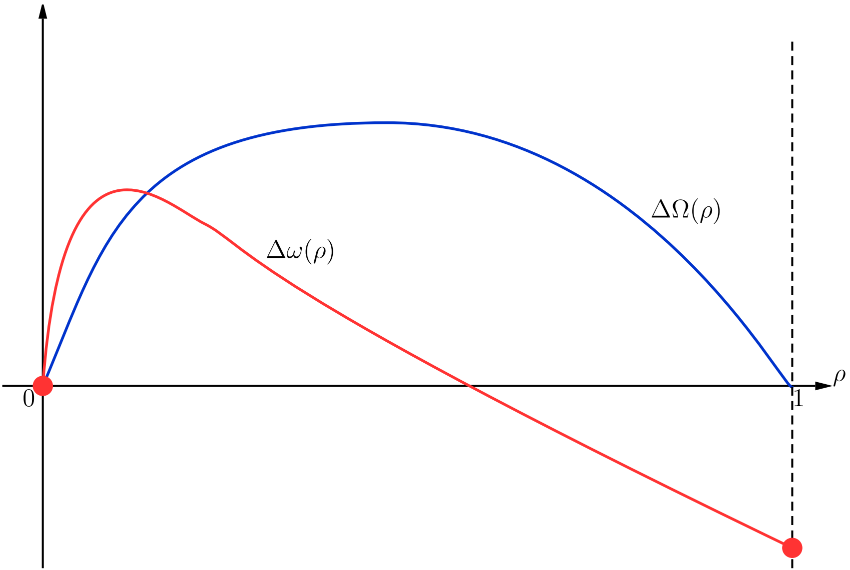

High Reputation ⇒ Stringent Rule with Separation

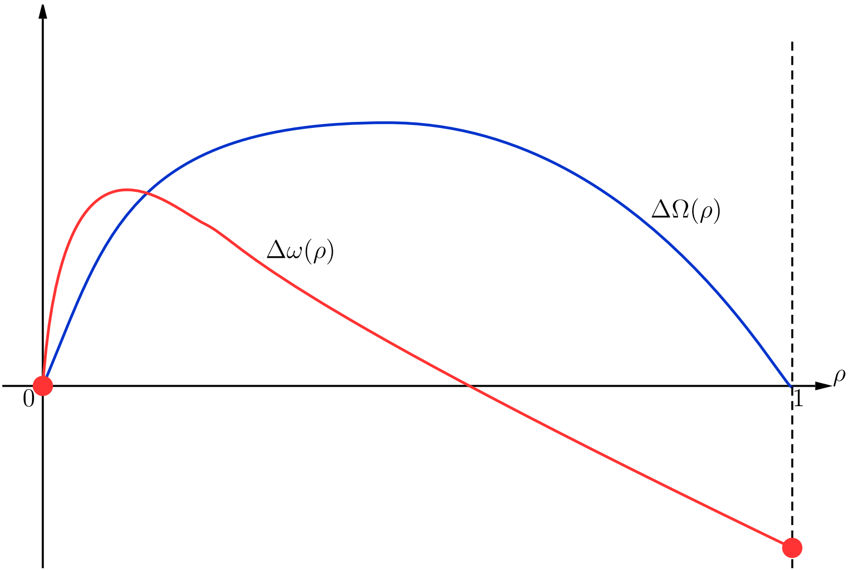

In Barro-Gordon model: Cutoff Rule

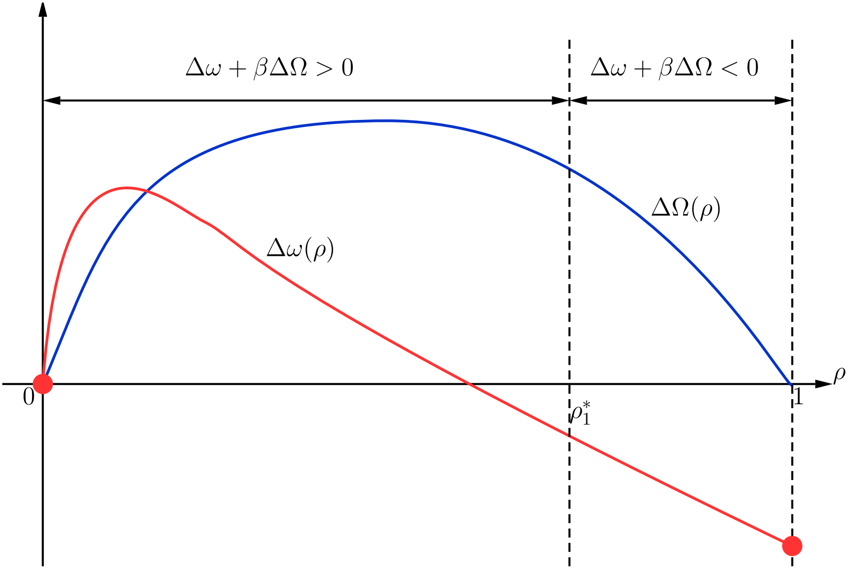

Optimal Dynamic Rule

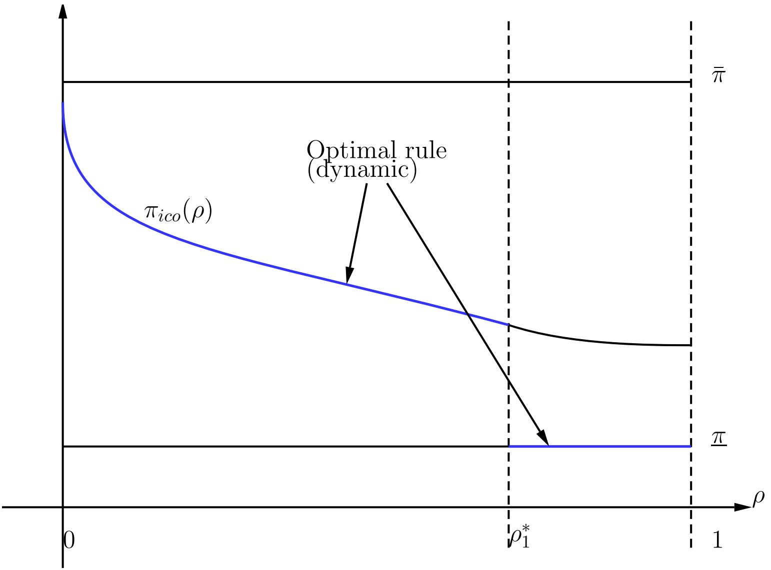

Result in the Context of our Two Examples

In the Bailout example

- $\pi_{ramsey} = \pi_0 = \underline{\pi} = 0$
  - Statically optimal rule is a strict no-bailout policy
  - Incentivizes maximal effort (minimal risk taking)
- Proposition 1: if reputation low allow for partial bailouts
  - On path-bailouts *necessary* to discipline future risk taking by banks

In the Barro-Gordon example

- $\pi_{ramsey} = \pi_0 = \underline{\pi} = 0$
  - Statically optimal rule is strict zero inflation target
- Proposition 1: if reputation low relaxed target is optimal

Extensions

- Insights from two period model extend to any finite horizon
  - Including the limit as $T \to \infty$ picture

- We also consider an extension in which the rules are "sticky"
  - Rule can be revised with probability $\alpha < 1$
  - As $T \to \infty$ main results unchanged

- Optimal rules when rule designer can commit
  - Rule designer also suffers from time inconsistency problem
  - There exists an interval of intermediate priors such that
    - Rule designer in period t wants to impose stringent rules in $t+1$
    - Rule designer in period $t+1$ chooses lenient rules

The Benefits of Opaque Rules

Role for Opaque Rules

When policy maker's type known:

- Transparent rules optimal because enable better monitoring
  - Provide incentives to policy-makers to not deviate
  - Avoid punishment on path

When policy makers's type uncertain:

- Opaque rules optimal because they help preserve uncertainty
  - Want rules that are hard to monitor (hard to detect deviation)
  - Complicated rules that have irrelevant contingencies

Optimal Degree of Monitoring

- Suppose private agents cannot observe π

- Observe a signal $\tilde{\pi} = \pi + \varepsilon$ where $\varepsilon \sim N\left(0, \sigma_\varepsilon^2\right)$

- Rule designer chooses σ_ε^2 as part of the optimal rule design
  - Rule is transparent if $\sigma_\varepsilon^2 = 0$
  - Rule is opaque if $\sigma_\varepsilon^2 > 0$

- The law of motion for beliefs is

$$\begin{aligned}\rho'(\tilde{\pi}, \rho) &= \frac{\rho \Pr(\tilde{\pi}|\pi_c)}{\rho \Pr(\tilde{\pi}|\pi_c) + (1-\rho)\Pr(\tilde{\pi}|\pi_o)} \\ &= \frac{\rho g(\tilde{\pi} - \pi_c|\sigma_\varepsilon)}{\rho g(\tilde{\pi} - \pi_c|\sigma_\varepsilon) + (1-\rho) g(\tilde{\pi} - \pi_o|\sigma_\varepsilon)}\end{aligned}$$

▸ Rule Designer's problem

Opaque Rules Are Optimal for High Reputation

Proposition

- *For* ρ *close to zero, it is optimal to pool and set* σ_ε *to zero;*
- *For* ρ *close to 1, it is optimal to separate and set* $\sigma_\varepsilon > 0$

Opaque Rules Are Optimal for High Reputation

Proposition

- *For* ρ *close to zero, it is optimal to pool and set* σ_ε *to zero;*
- *For* ρ *close to 1, it is optimal to separate and set* $\sigma_\varepsilon > 0$

- Low $\rho \Rightarrow$ optimal to pool:
  - $\sigma_\varepsilon = 0$ to relax IC for optimizing type
  - Same logic as in ACK
- High $\rho \Rightarrow$ optimal to separate
  - Spse for contradiction $\sigma_\varepsilon = 0 \Rightarrow \rho' \in \{0, 1\}$ and $\pi_o = \pi^*(x)$
  - Can support same policies by choosing $\sigma_\varepsilon = \infty \Rightarrow \rho' = \rho$
  - Since $W(\rho) > \rho W(1) + (1 - \rho) W(0)$ have an improvement. Contradiction.

Opaque Rules Are Optimal for High Reputation

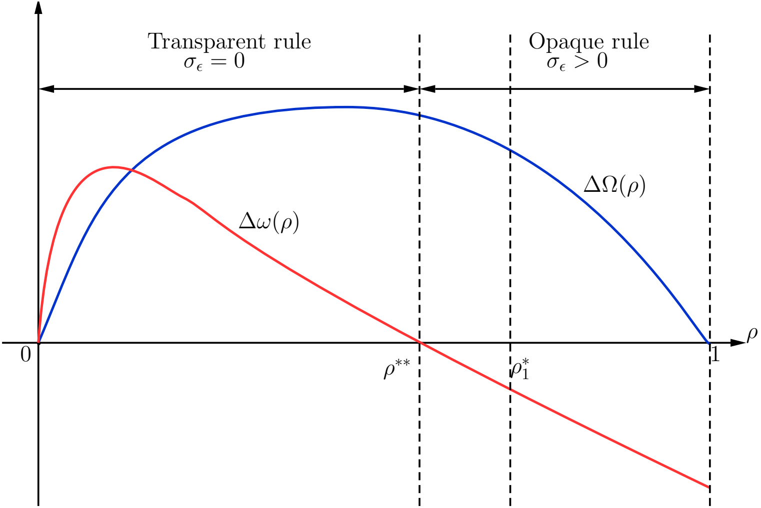

With opaque rules no trade-off b/w dynamic and static benefits

Optimal Tenure

- Replacing policy maker:
  - Get the static benefits of separation without the dynamic losses

- Equivalent to choosing a perfectly opaque rule with $\sigma_\varepsilon = \infty$

- Rule designer's problem as before w/ restriction $\sigma_\varepsilon \in \{0, \infty\}$

Proposition

In the Barro-Gordon model, there exists $\rho^{**} < \rho_1^*$ such that:

- *For $\rho \leqslant \rho^{**}$ do not terminate and pool*
- *For $\rho \geqslant \rho^{**}$ terminate and separate*

Random Rules are Optimal for High Reputation

- Allow for randomization in π_c
  - Make rules contingent on irrelevant details

Proposition

In our two examples:

- *For ρ close to 0, a deterministic rule is optimal*
  - $\pi_c = \pi_{ico}(\rho)$ *with probability one*
- *For ρ close to 1, it is optimal to have stochastic rules.*

Random Rules are Optimal for High Reputation, cont.

- Consider ρ close to 1
- Spse optimal to separate and no randomization
- Consider the perturbation

$$\pi_c^{dev} = \begin{cases} \pi_c & \text{with pr } 1-\varepsilon \\ \pi_o & \text{with pr } \varepsilon \end{cases} \Rightarrow \rho' = \begin{cases} \frac{\rho\varepsilon}{\rho\varepsilon+(1-\rho)} > 0 & \text{if } \pi = \pi_o \\ 1 & \text{if } \pi = \pi_c \end{cases}$$

  with value $W^{dev}(\varepsilon) - W \approx [(1-\beta)\Delta\omega'(\varepsilon) + \beta\Delta\Omega'(\varepsilon)]\,\varepsilon$

- Since $\lim_{\rho\to 1}\lim_{\varepsilon\to 0}\Delta\Omega'(\varepsilon) \to \infty$ while $|\Delta\omega'(\varepsilon)| < M$
  $\Rightarrow$ the perturbation is profitable
- Key inputs
  - $W_0' > 0$
  - $\lim_{\rho\to 1}\lim_{\varepsilon\to 0}\partial\rho'/\partial\varepsilon = \infty$

Conclusion

- Optimal design of rules when uncertain about whether policy maker follows the rule ex-post

- If reputation low optimal to design lenient rules which help preserve uncertainty

- Opaque/complicated rules desirable as they preserve reputation without the static costs of leniency

ADDITIONAL SLIDES

Example 1: Barro-Gordon

- π: inflation rate

- $\tilde{x}$: wage inflation set by unions as

$$\tilde{x} = \phi(\mathbb{E}\pi) = \mathbb{E}\pi$$

  and $x = -\tilde{x}$

- Preferences

$$w(x, \pi) = -\frac{1}{2}\left[(\psi - x - \pi)^2 + \pi^2\right]$$

  with inflation bias $\psi > 0$

Back

Example 2: Bank-Bailout a la Kareken-Wallace

- Private agents: bank, lenders
- Bank:
  - Raise 1 to finance an investment opportunity
  - Promises to repay R but there is limited liability
  - Chooses effort e, disutility $\nu(e)$ w/ $\nu' > 0$, $\nu'' > 0$
- Returns from investment are
  - R_H with probability $p(e)$ w/ $p' > 0$, $p'' < 0$
  - 0 with probability $1 - p(e)$
- Social cost of default given by $\psi(1 - \pi)$
- Policy maker can avoid defaults w/ transfer to bank/lenders
  - $\pi \in [0, 1]$: recovery rate after a bad realization
  - Taxation cost associated with transfers $c(\pi)$, $c' \geqslant 0, c'' \geqslant 0$

Example 2: Bank-Bailout a la Kareken-Wallace, cont.

- π: recovery rate

- $x = e$,

$$\phi(\mathbb{E}\pi) = \arg\max_{e} -v(e) + p(e)(R_H - R(e))$$

subject to

$$R(e) = \frac{1 - (1 - p(e))\mathbb{E}\pi}{p(e)}$$

- Social welfare function

$$w(x, \pi) = -v(x) + p(x)R_H + (1 - p(x))\psi(1 - \pi) - c(\pi)$$

◂ Back

Infinite Horizon

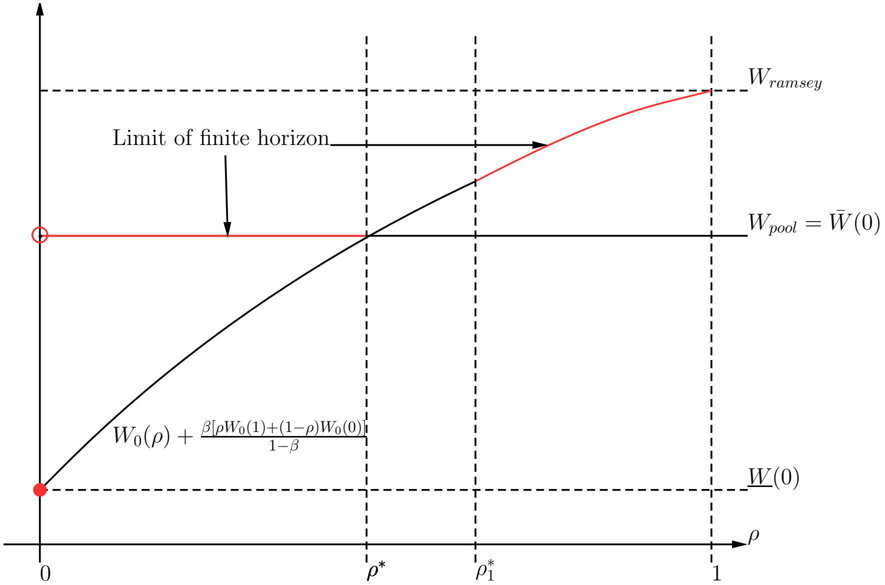

Back

Rule Designer's Problem

$$\max_{x,\pi_c,\pi_o,\sigma_\varepsilon} \rho \left[w(x,\pi_c) + \beta \int W_0\left(\rho'(\pi_c+\varepsilon,\rho)\right) g(\varepsilon|\sigma_\varepsilon)\, d\varepsilon \right]$$
$$+ (1-\rho)\left[w(x,\pi_o) + \beta \int W_0\left(\rho'(\pi_o+\varepsilon,\rho)\right) g(\varepsilon|\sigma_\varepsilon)\, d\varepsilon \right]$$

subject to

$$x = \phi\left(\rho, \rho\pi_c + (1-\rho)\pi_o\right),$$

$$w(x,\pi_o) + \beta_o \int V_0\left(\rho'(\pi_o+\varepsilon,\rho)\right) g(\varepsilon|\sigma_\varepsilon)\, d\varepsilon \geqslant$$
$$\geqslant w(x,\pi) + \beta_o \int V_0\left(\rho'(\pi+\varepsilon,\rho)\right) g(\varepsilon|\sigma_\varepsilon)\, d\varepsilon \quad \forall \pi$$

and

$$\rho'(\tilde{\pi},\rho) = \frac{\rho g(\tilde{\pi}-\pi_c|\sigma_\varepsilon)}{\rho g(\tilde{\pi}-\pi_c|\sigma_\varepsilon) + (1-\rho) g(\tilde{\pi}-\pi_o|\sigma_\varepsilon)}$$

Back