Robustly Optimal Monetary Policy with Near-Rational Expectations

Michael Woodford

Columbia University

NBER Summer Institute, July 2006

- An extensive literature considers optimal monetary policy under the assumption of *rational expectations*

- Key conclusion: important to take account of effects of *commitment of future policy* on PS *expectations*

- Implications:
  - optimal commitment $\neq$ discretionary policy
  - optimal policy should be *history-dependent*

- Example: optimal policy problem in CGG (1999):

$$\min \quad E_0 \sum_{t=0}^{\infty} \beta^t [\pi_t^2 + \lambda (x_t - x^*)^2]$$

$$\text{s.t.} \qquad \pi_t = \kappa x_t + \beta E_t \pi_{t+1} + u_t$$

- optimal commitment $\Rightarrow$ long-run average inflation $\pi^* = 0$

  — whereas discretion results in $\pi > 0$ on average (if $x^* > 0$)

- optimal commitment $\Rightarrow$ should subsequently *undo* price changes due to cost-push shocks

  — whereas discretion would “let bygones be bygones”

- Question: isn't it unrealistic to assume that any pattern created by systematic policy must be *correctly anticipated* by the private sector?

- Perhaps such an assumption exaggerates the importance of *commitment* and *history-dependence*

- An alternative: assume a particular *model* of how expectations are formed from past observations

— Example: Orphanides and Williams (2005): *least-squares learning* $\Rightarrow$ a *different* output/inflation variability tradeoff than CB would expect if assumes RE

— More aggressive inflation control in face of cost-push shocks is optimal, within a simple family of linear policies

— Gaspar, Smets and Vestin (2005, 2006): dynamic optimization by a CB that recognizes that PS expectations result from least-squares learning

- Objection: choosing optimal policy for a *specific model* of expectations may *exaggerate* CB's ability to model PS behavior, lead to overly ambitious attempts to manipulate expectations

  — the pitfall that RE were introduced to avoid!

- Another alternative: let CB recognize that PS expectations may *differ* from the probabilities implied by its own model

  — *not* assume that it *knows* what PS expectations must be, in case of a particular policy rule

  — might be *any* beliefs, among those not *too different* from what CB's model implies ("NRE")

  — choose the policy that is *least vulnerable* to deviation of PS expectations from model-consistency

  — as in theories of "ambiguity aversion," "robust control"

- Not a standard “robust control” problem:
  - control of a forward-looking system: advance commitment matters
  - no uncertainty on part of policymaker about correct *model* (given PS expectations)
- Also different from robust control problem for forward-looking system considered by Hansen and Sargent (2005, chap. 16):
  - failure of PS to optimize under RE not due to its uncertainty about correct model
  - the distorted beliefs of PS that matter are those that would be most inconvenient *for the policymaker,* not those that would be most *to private agents.*

Application

Policy problem considered under RE by CGG (1999)

- Model: NK Phillips curve w/'cost-push' shock

$$\pi_t = \kappa x_t + \beta \hat{E}_t \pi_{t+1} + u_t$$

— $\kappa > 0, \quad 0 < \beta < 1, \quad u_t$ exogenous

- $\hat{E}_t[\cdot]$: expectation wrt *common* distorted beliefs of all private agents

— usual microfoundations of NKPC still apply

- CB objective: minimize

$$L_0 \equiv E_0 \sum_{t=0} \beta^t [\pi_t^2 + \lambda (x_t - x^*)^2]$$

where $\lambda > 0, \quad x^* \geq 0.$

- $E_t[\cdot]$: expectation under CB beliefs, treated as *true*
- Law of motion of exogenous state (under CB beliefs):

$$s_{t+1} = A s_t + B w_{t+1}, \qquad w_{t+1} \sim N(0, I), \quad \text{i.i.d.}$$

$$u_t = v' s_t$$

- CB strategy: commits itself (at initial date) to a state-contingent policy

$$\pi_t = \pi(h_t)$$

where $h_t \equiv (s_t, s_{t-1}, \ldots)$, *without knowledge* of what distorted private beliefs will be

- Assume no problem *implementing* the policy: can observe private expectations then

- Any such strategy uniquely defines evolution $\{\pi_t, x_t\}$ (*given* PS beliefs), and hence CB expected losses

- This choice of strategy space *not* innocuous (as in CGG)!

Near-Rational Expectations

- PS beliefs must be *absolutely continuous* wrt truth [over any finite time interval]

$\Rightarrow$ there exists a process $\{m_{t+1}\}$ with

$$m_{t+1} \geq 0 \text{ a.s.}, \qquad E_t[m_{t+1}] = 1.$$

such that

$$\hat{E}_t[X_{t+1}] = E_t[m_{t+1} X_{t+1}]$$

for any random X_{t+1}

- Degree of distortion of PS beliefs measured by *relative entropy*

$$R_0 \equiv E_0 \sum_{t=0}^{\infty} \beta^t m_{t+1} \log m_{t+1}$$

as in Hansen-Sargent treatment of robust policy

  - a positive-valued, convex function of distorted prob. measure, uniquely minimized (= 0) when $m_{t+1} = 1$ a.s. [case of RE]
  - a measure of how easily the distorted beliefs should be disconfirmed by data [according to CB beliefs]
  - discounting at rate β means CB concern with potential PS misunderstanding doesn't vanish asymptotically

- Two possible versions of the "NRE" constraint:
  - CB fears PS beliefs m that are worst for its objective L_0, subject to constraint

    $$R_0 \leq \bar{R},$$

    for some $0 < \bar{R} < +\infty$ [Gilboa-Schmeidler]

  - CB fears PS beliefs m that maximize

    $$L_0 - \theta R_0$$

    for some $\theta > 0$ [Hansen-Sargent; Maccheroni *et al.*]

- Here use the "multiplier" form

- Worst-case NRE: given CB policy $\{\pi_t\}$, PS beliefs solve

$$\max_{\{m_{t+1}\}} \quad E_0 \sum_{t=0} \beta^t [\pi_t^2 + \lambda (x_t - x^*)^2]$$
$$-\theta E_0 \sum_{t=0}^{\infty} \beta^t m_{t+1} \log m_{t+1}$$

s.t. $\pi_t = \kappa x_t + \beta E_t [m_{t+1} \pi_{t+1}] + u_t$ for each t

$E_t m_{t+1} = 1$ for each t

- CB chooses $\{\pi_t\}$ to min this max'd value

Robustly Optimal Linear Policy

- Here further restrict CB strategy space to *conditionally linear* policies of form

$$\pi_{t+1}(w_{t+1}) = p_t^0 + p_t^{1\prime} w_{t+1}$$

for some p_t^0, p_t^1 depending only on h_t.

- Note: optimal policy is of this form in case of RE (and p_t^1 is not time-varying)

- Worst-case NRE (1): interior solution exists only if

$$|p_t^1|^2 < \frac{\theta}{\beta^2}\frac{\kappa^2}{\lambda} \qquad (*)$$

— otherwise, objective of "perverse agent" is *convex,* and worst-case results in *unbounded losses* for CB

- Obviously, CB should choose policy that satisfies (*) at all times

— contrasts with result under RE: optimal p^1 proportional to σ_u

— at least for large σ_u, evident that concern for robustness leads to *less sensitivity* of inflation to cost-push disturbances

- Worst-case NRE (2): in the case of a linear policy (p_t^0, p_t^1) satisfying (*), CB fears that PS will expect w_{t+1} to be conditionally distributed as $N(\mu_t, I)$, where
  - if $p_t^1 = 0$, $\mu_t = 0$, but
  - if $p_t^1 \neq 0$,

$$\mu_t = (\bar{\pi}_t - p_t^0)\ \frac{p_t^1}{|p_t^1|^2},$$

where the worst-case inflation expectation is

$$\bar{\pi}_t = \Delta_t^{-1}\left[p_t^0 - (\pi_t - u_t - \kappa x^*)\frac{\beta\lambda}{\theta\kappa^2}|p_t^1|^2\right],$$

$$\Delta_t \equiv 1 - \frac{\beta^2}{\theta}\frac{\lambda}{\kappa^2}|p_t^1|^2 > 0$$

- Intuition: worst-case NRE distort $\hat{E}_t \pi_{t+1}$ away from p_t^0 in the direction *opposite* to that needed to bring x_t closer to x^*

  — distortion is greater the larger is $|p_t^1|$, becoming unboundedly large as the bound (*) is approached

- CB therefore fears output gap

$$x_t^{pess} - x^* = \frac{(\pi_t - u_t - \kappa x^*) - \beta p_t^0}{\kappa \Delta_t}$$

  — larger by factor $\Delta_t^{-1} \geq 1$ than under RE

- CB's Robust Conditionally Linear Policy Problem: choose $\{p_t\} \equiv \{p_t^0, p_t^1\}$ to

$$\min E_0 \sum_{t=0}^{\infty} \beta^t L(p_{t-1}; p_t; s_t)$$

where

$$\begin{aligned} L(p_{-1}; p; s) \equiv & \quad \frac{1}{2}(p_{-1}^0 + p_{-1}^{1\prime} w)^2 \\ + & \quad \frac{1}{2}\frac{\lambda}{\kappa^2}\frac{[(p_{-1}^0 + p_{-1}^{1\prime} w) - u - \kappa x^* - \beta p^0]^2}{\Delta(p^1)} \end{aligned}$$

- Solution to this problem is not a *linear* policy (= inflation a linear function of h_t), except in case of RE

- Useful to consider optimal policy within a more restrictive class: assume that $\{p_t^1\}$ is a *deterministic* sequence, chosen to minimize *unconditional* expected losses

  — again a restriction satisfied by optimal policy under RE

  — Optimal policy then makes inflation a *linear* function of history of states

- Optimal policy from a *timeless perspective:* (a) choose $\{p_t\}$ for $t \geq 0$ to

$$\min E \sum_{t=0}^{\infty} \beta^t L(p_{t-1}; p_t; s_t)$$

given p_{-1};

- (b) *Self-consistent* choice of p_{-1}:
  - $p_{-1}^1 = \bar{p}^1$, the constant value that is then optimal for $t \geq 0$;
  - (p_{-1}^0, s_0) drawn from joint distribution equal to the ergodic distribution for (p_{t-1}^0, s_t) under the optimal dynamics for $t \geq 0$.

- Call $\{p_t^0\}, \bar{p}^1$ that solve this problem the *robust linear policy.*

Characterization

- Robust linear policy (1): consider the sub-problem: choose (robustly) optimal $\{p_t^0\}$ given $p_t^1 = \bar{p}^1$ (arbitrary vector) for all t

- This is an LQ problem with constant coefficients: optimal policy a *linear* function of the history of shocks

- Only difference in this problem due to concern for robustness: relative weight on second term in loss function

$$\frac{\lambda}{\kappa^2} \quad \rightarrow \quad \frac{\lambda}{\kappa^2 \Delta}$$

is *increased* ($\theta^{-1} > 0 \quad \Rightarrow \quad \Delta < 1$)

- Solution:

$$(1-\mu L)p_t^0 = -\mu E_t[(1-\beta\mu L^{-1})^{-1}(1-L^{-1})u_t]+\mu\bar{p}^{1\prime}w_t \tag{1}$$

where $0 < \mu < 1$ is the smaller root of

$$\beta\mu^2 - \left(1+\beta+\frac{\kappa^2\Delta}{\lambda}\right)\mu + 1 = 0$$

- Optimal long-run *average* inflation target is *zero,* regardless of size of x^*, even when $\theta^{-1} > 0$

  — *no* “inflation bias”!

- Example: i.i.d. cost-push shocks: w_t scalar, $u_t = \sigma_u w_t$

$$p_t^0 = \mu p_{t-1}^0 - \mu(\sigma_u - \bar{p}^1) w_t$$

- As long as $\bar{p}^1 < \sigma_u$ (shown below), still optimal to commit to *lower* subsequent inflation in response to a positive cost-push shock, as under RE

- Note that concern for robustness *lowers* Δ $\Rightarrow$ raises μ (closer to 1) $\Rightarrow$ greater intrinsic persistence in average inflation target

  — greater *history-dependence* is optimal

  — for given $\bar{p}^1$, it also *increases* the optimal response of subsequent inflation to the shock, both initially, and (even more) cumulatively

- Robust linear policy (2): FOCs for optimal choice of sequence $\{p_t^1\}$:

$$E[L_{p^1}(p_{t-1}; p_t; s_t) + \beta L_{p^1_{-1}}(p_t; p_{t+1}; s_{t+1})] = 0$$

— we seek steady state $\bar{p}^1$, implicitly defined by

$$E[L_{p^1}(p^0_{t-1}, \bar{p}^1; p^0_t, \bar{p}^1; s_t) + \beta L_{p^1_{-1}}(p^0_t, \bar{p}^1; p^0_{t+1}, \bar{p}^1; s_{t+1})] = 0 \tag{2}$$

- Robust linear policy is then a pair $\{p_t^0\}$, $\bar{p}^1$, such that $\{p_t^0\}$ is determined by (1) given $\bar{p}^1$, and $\bar{p}^1$ satisfies (2) given $\{p_t^0\}$.

- Optimal policy under RE is a linear policy of same form:

  — $\{p_t^0\}$ determined by (1) given $\bar{p}^1$; but μ is root corresponding to $\Delta = 1$;

  — first term vanishes in (2); in i.i.d. case,

$$\bar{p}^1 = \mu\sigma_u.$$

- Consequences of concern for robustness:

  — $\Delta < 1$ implies larger value of $\mu < 1$;

  — in i.i.d. case, additional term in FOC (2) implies

$$\bar{p}^1 < \mu\sigma_u$$

  — may be *much* lower in case of large σ_u (and small θ), as required by bound (*).

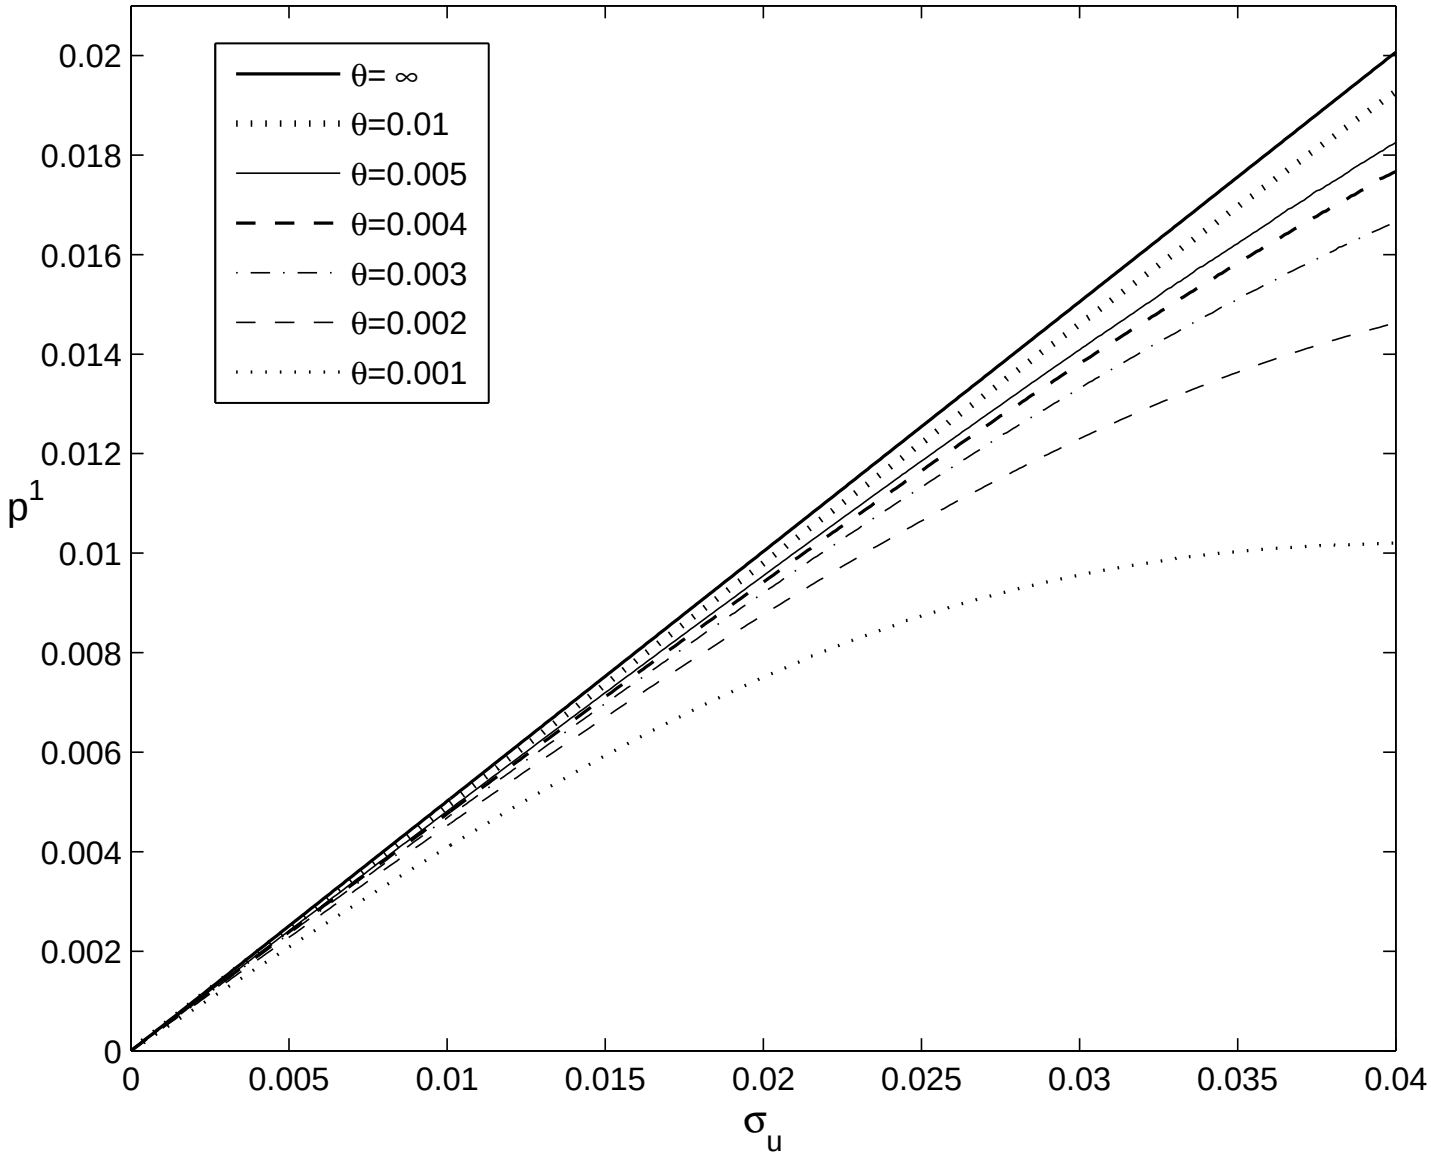

Figure 1: Variation of $\bar{p}^1$ with σ_u, under alternative degrees of concern for robustness.

- Impulse responses to a cost-push shock, under robust linear policy [i.i.d. case]:

  — inflation *increased* by positive shock ($\bar{p}^1 > 0$), but often by less than would be true under RE

  — CB plans *lower* inflation in subsequent periods

  — cumulative *disinflation* in periods following shock

$$-\frac{\mu}{1-\mu}(\sigma_u - \bar{p}^1)$$

  more than offsets the initial inflation ($\bar{p}^1$) if $\bar{p}^1 < \mu\sigma_u$, which is always the case if $\theta^{-1} > 0$.

- Because PS inflation expectations fall by *less* than the amount of disinflation planned by the CB, CB should commit to *even more* subsequent disinflation than otherwise necessary!

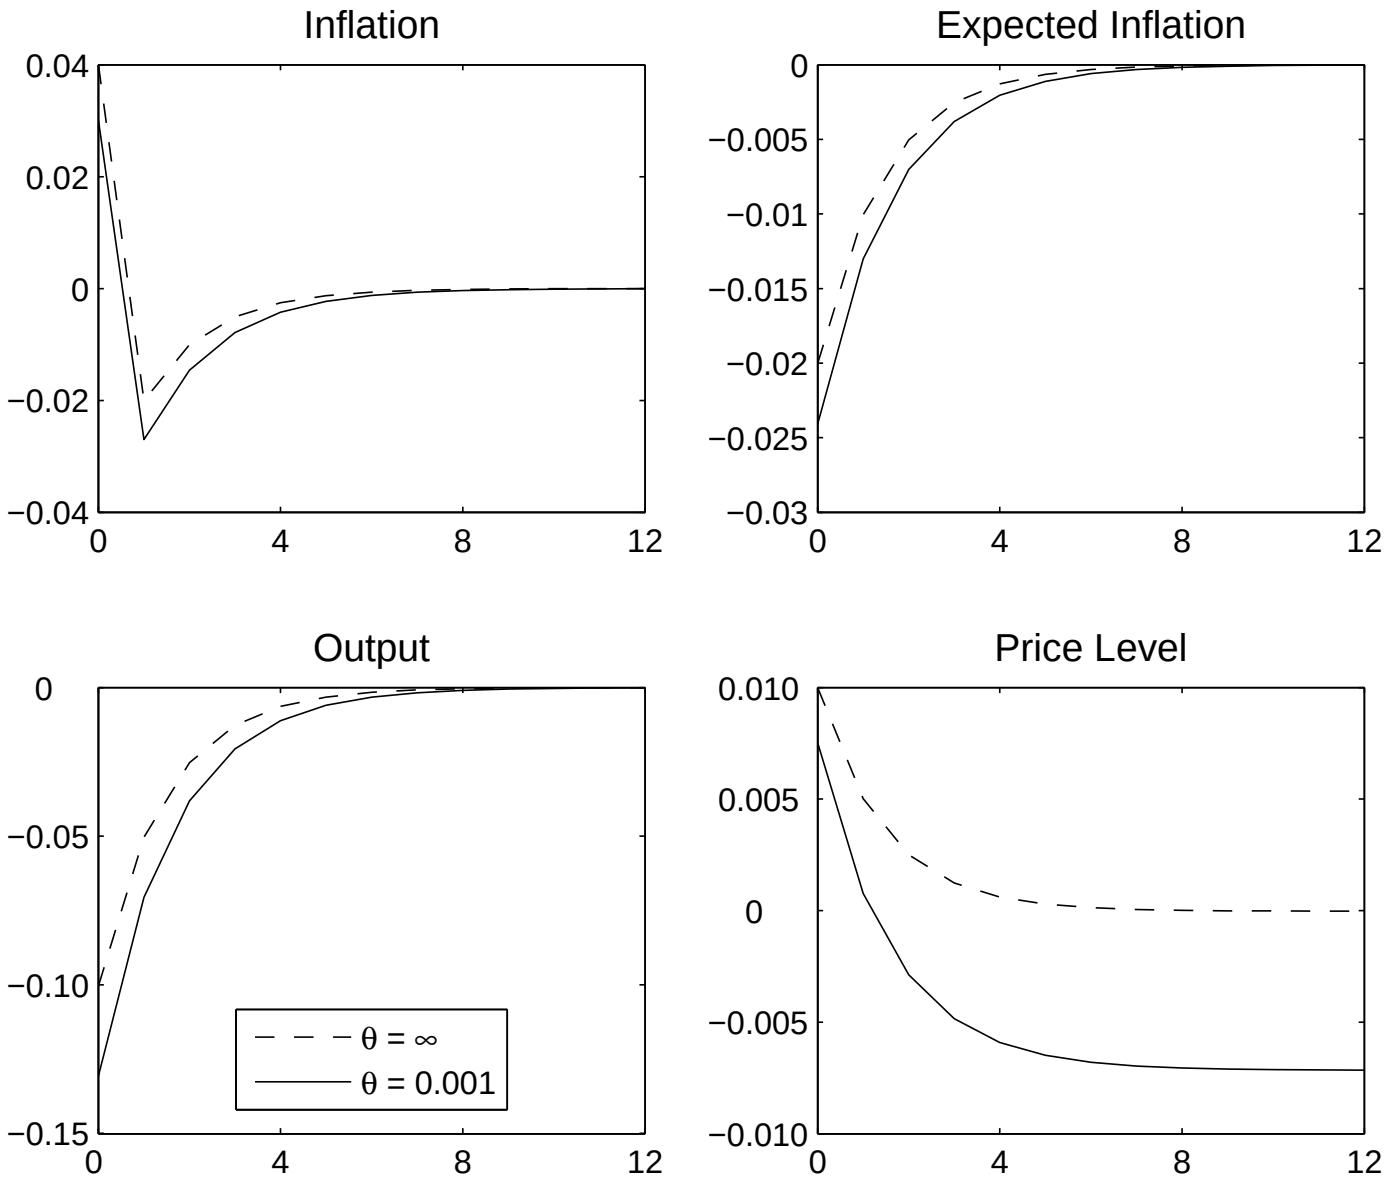

Figure 2: Optimal responses to a positive cost-push shock, with and without concern for robustness.

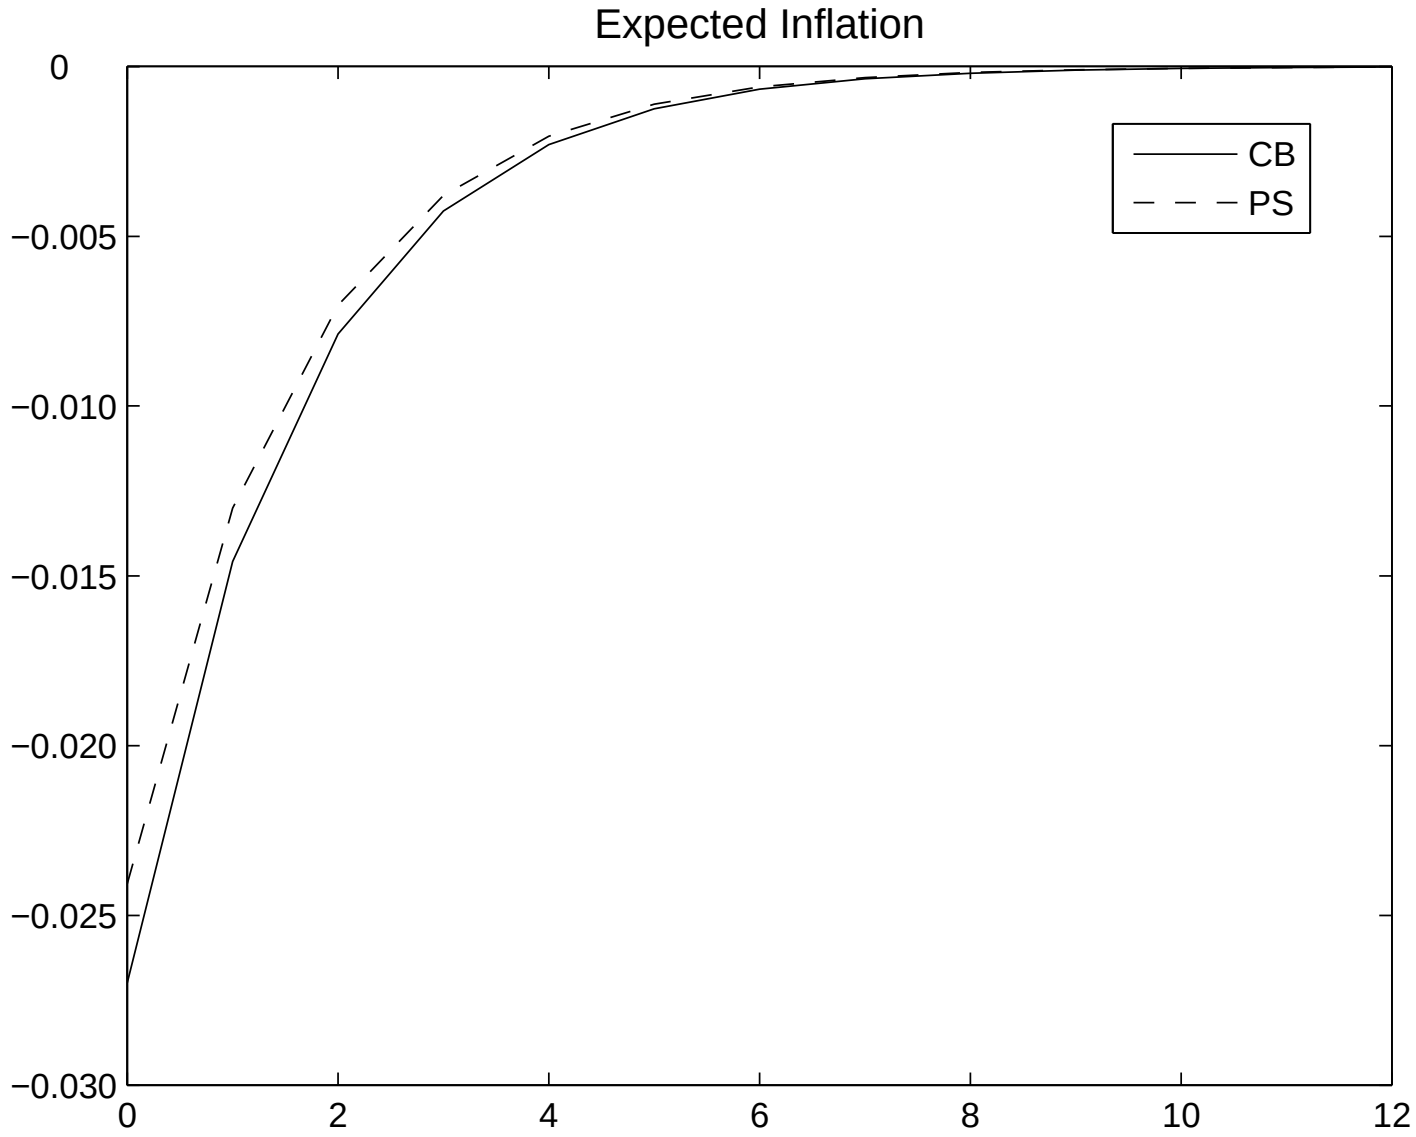

Figure 3: Distortion of PS beliefs, in the worst-case NRE contemplated by the CB when θ= 0.001.

Optimizing Policy under Discretion

- How different is optimal commitment from what an optimizing CB would choose under discretion?

- Suppose each period CB chooses π_t after observing s_t, making no commitment regarding inflation in subsequent periods

  — in a Markov eq'm, π_t will be a time-invariant function of s_t

  — in choosing π_t, it is assumed that the function $\pi(\cdot)$ will describe CB policy in all later periods

- Optimization by CB under discretion implies that for any state s_t, $\bar{\pi}(s_t)$ solves the problem

$$\min_{\pi_t} V(\pi_t; s_t), \tag{3}$$

where $V(\pi_t; s_t)$ is expected CB loss from t onward, given choice of π_t and expected policy $\bar{\pi}$ in all later periods.

- Given expected CB behavior of this kind, *worst-case NRE beliefs* solve

$$\max_{m_{t+1}(\cdot)} \frac{1}{2}[\pi_t^2 + \lambda(x_t - x^*)^2] - \theta E_t \beta^t m_{t+1} \log m_{t+1}$$

$$+ \beta E_t V(\bar{\pi}(s_{t+1}); s_{t+1}), \qquad (4)$$

where x_t satisfies

$$\pi_t = \kappa x_t + \beta E_t[m_{t+1}\bar{\pi}(s_{t+1})] + u_t$$

- A *robust MPE* is then a pair of functions $\bar{\pi}(\cdot)$ and $V(\cdot;\cdot)$ such that

  [i] for any pair $(\pi_t; s_t)$, $V(\pi_t; s_t)$ is the maximized value of (4); and

  [ii] for any state s_t, $\bar{\pi}(s_t)$ is the inflation rate that solves the problem (3).

- A *robust linear MPE* is a robust MPE in which $\bar{\pi}(\cdot)$ is a linear function of the state,

$$\bar{\pi}(s_t) = \pi^* + d' s_t$$

- Previous characterization of worst-case beliefs still applies
- Hence $\bar{\pi}(s_t)$ solves

$$\min_{\pi_t} L(\pi_t; \bar{p}(s_t); s_t)$$

- Case of i.i.d. shocks: this has one, two, or zero linear solutions, depending on the size of σ_u

  — when two solutions, MPE on lower branch are "expectationally stable"

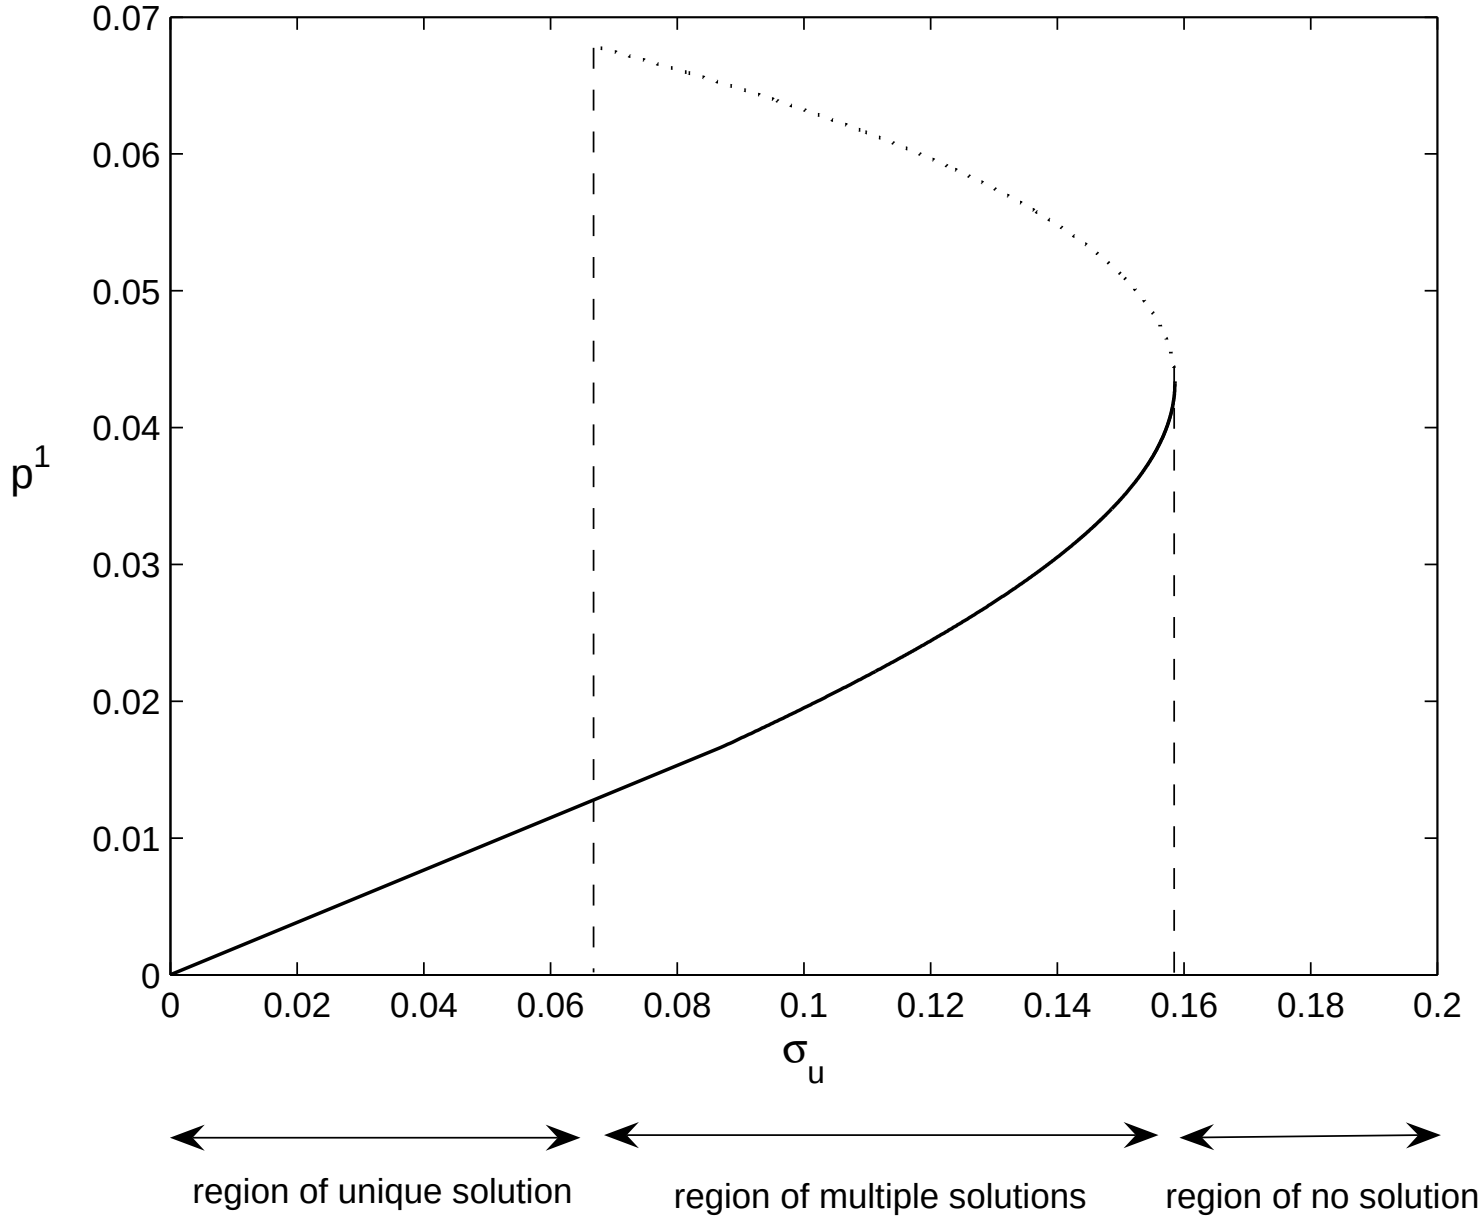

Figure 4: Varying numbers of linear MPE, depending on size of cost-push shocks.

- Average inflation rate in robust linear MPE:

$$\pi^* = \frac{\lambda \kappa x^*}{\kappa^2 \bar{\Delta} + (1-\beta)\lambda} > 0$$

— higher the smaller is $\bar{\Delta}$, so inflation bias is *increased* by concern for robustness

- Response to cost-push shocks in robust linear MPE [i.i.d. case]:

$$\bar{p}^1 = \frac{\lambda}{\kappa^2 \bar{\Delta} + \lambda} \sigma_u$$

— as concern for robustness lowers $\bar{\Delta}$, $\bar{p}^1$ is *higher* than under RE

— thus concern for robustness moves $\bar{p}^1$ under discretion *farther* from the value associated with optimal linear commitment

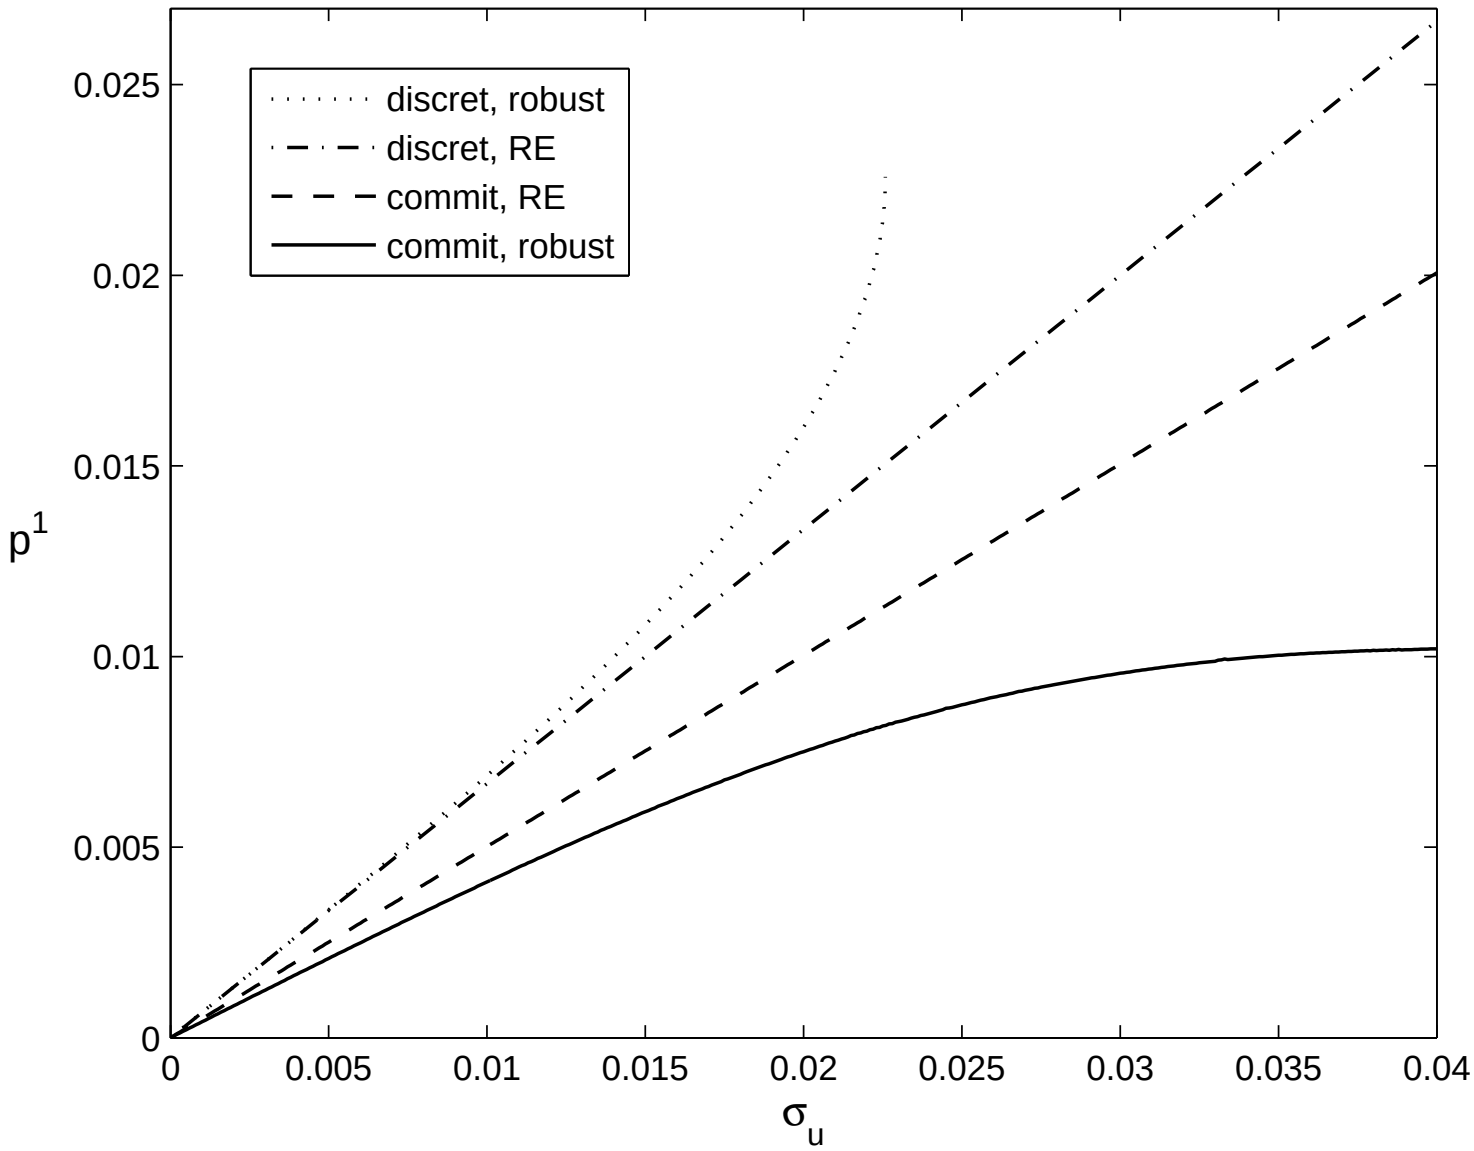

Figure 5: Variation of $\bar{p}^1$ with σ_u, under discretionary policy and under an optimal commitment.

Conclusions

- It is possible to analyze optimal policy for a central bank that recognizes that PS expectations may not be model-consistent, without committing oneself to a *particular* model of expectational error

- Even when the CB's uncertainty about PS expectations is considerable (the case of low θ), calculation of the effects of PS *anticipations* of the systematic component of policy is still quite important

  — *commitment* still important

  — optimal policy still *history-dependent*

- It may well be optimal for CB to allow *less* immediate impact of cost-push shocks on inflation, when cannot be certain PS will correctly anticipate policy, as argued by Orphanides and Williams

  — yet at the same time, it is typically desirable to *increase* the sensitivity of inflation to *past* shocks