A Decision-Theoretic Approach to Monetary Policy under Weak Identification

Jorge de la Cal Medina

University of Amsterdam, Tinbergen Institute
European Central Bank, Directorate General Economics

IAAE 2026

Disclaimer

The views expressed in this presentation are those of the author and do not necessarily reflect those of the European Central Bank or the Eurosystem.

Introduction

- Common practice: evaluate policy in DSGE model at point estimate (Schmitt-Grohé and Uribe, 2007; Galí, 2015)
- But DSGE models may be observationally equivalent
  - here focus on impulse responses (Christiano et al., 2005)
  - different structural parameters can fit the same moments
- The parameter-to-moment mapping may be many-to-one
  - Chari et al. (2009): reduced-form fit can mask different structures
  - Mavroeidis et al. (2014): NKPC slope can be weakly identified
  - Canova and Sala (2009): IRFs may leave structural parameters weakly identified
- Implication: plug-in policy can be misleading

This Paper

Question: how should a policymaker choose one rule when welfare-relevant parameters are (potentially) weakly identified?

- Treat policy choice as a statistical decision problem
  - action: operational interest-rate rule
  - state: structural parameter
  - loss: welfare cost of the rule
- Apply "optimal decision rules for weak gmm" (Andrews and Mikusheva, 2022)
  - use a Bayes rule
  - (quasi-) posterior makes trade-off between structural values explicit
  - choose the rule with lowest posterior expected loss
- Show policy consequences
  - Monte Carlo: regret relative to the oracle policy rule
  - medium-scale NK model: rule disagreement on US data

Relation to Robust Policy

- Hansen–Sargent robustness (Hansen and Sargent, 2008)
  - problem: misspecification around a benchmark model
  - response: choose against least-favorable distortions
- Giannoni–Woodford robustness (Giannoni and Woodford, 2005, 2017)
  - problem: optimal policy depends on uncertain model dynamics
  - response: target criteria that remain valid across a class of model dynamics
- This paper
  - problem: weak identification of welfare-relevant parameters
  - response: choose a rule by posterior expected welfare

Structure of Presentation

Policy Problem and Weak Identification

Decision-Theoretic Response

Empirical Evidence

Conclusion

The Policy Problem

- If ψ were known, the policymaker would solve

$$a^\star(\psi) \in \arg\max_{a \in \mathcal{A}} W(a, \psi).$$

  - a: interest-rate policy
  - ψ: policy-invariant structural parameters
  - $W(a, \psi)$: welfare criterion

- The plug-in policy replaces ψ by an estimate

$$\widehat{a}^{\mathrm{PI}} = a^\star(\hat{\psi}).$$

- **Problem:** plug-in evaluates policies at one point
  - disregards sampling uncertainty
  - severe when ψ is weakly identified

Impulse-Response Matching

- Empirical DSGE work often estimates θ by matching impulse responses
  - Christiano et al. (2005) discipline a medium-scale DSGE model with IRFs
  - Hall et al. (2012) methodology formalization
- Minimum-distance estimator: $\hat{\theta}^{MD} = \arg\min_{\theta\in\Theta} Q_T(\theta)$,

$$Q_T(\theta) = \{\hat{\gamma}_T - g(\theta)\}' \widehat{V}^{-1} \{\hat{\gamma}_T - g(\theta)\}$$

  - $\hat{\gamma}_T$: empirical impulse responses
  - $g(\theta)$: model-implied impulse responses
  - $\widehat{V}$: weighting matrix
- Strong identification
  - the population criterion is uniquely minimized at θ_0
- Weak identification
  - local to identification failure

Policy as a Statistical Decision Problem

- State: $\theta = (\psi, \alpha)$
  - ψ: policy-invariant structural parameters
  - α: sample policy environment
- Data: empirical impulse responses $\hat{\gamma}_T$
- Action: interest-rate policy $a \in \mathcal{A}$
- Loss: $L(a, \psi) = -W(a, \psi)$
- Decision rule: $\delta : \mathcal{X} \to \mathcal{A}$ maps data to policy
- Criterion: choose a rule with low risk (expected loss)

$$\mathbb{E}_{\theta_0}\left[L\{\delta(X_T), \psi\}\right]$$

Optimality Criterion

- Need an optimality criterion to choose a decision rule
  - no uniformly best rule under weak identification
- Admissibility
  - rules out dominated decision rules
  - many admissible rules can remain
- Minimax and minimax regret
  - hedge against the worst-case state
  - can be overly conservative
  - not guaranteed to select unique rule
- Bayes risk
  - averages risk across plausible structural values
  - implies admissibility
  - makes the trade-off across states explicit
  - under strict curvature, it delivers a unique rule

Quasi-Bayes for IRF Matching

- The Bayes policy minimizes posterior expected loss
$$\widehat{a}^{\mathrm{QB}} = \arg\min_{a \in \mathcal{A}} \int L(a, \psi)\, \pi(\theta \mid X_T)\, d\theta.$$
- For IRF matching, use the quasi-posterior
$$\pi(\theta \mid X_T) \propto \exp\left[-\frac{T}{2} Q_T(\theta)\right] \pi(\theta).$$
- Based on the MD criterion (Chernozhukov and Hong, 2003; Andrews and Mikusheva, 2022)
- Not a full posterior: no likelihood for the whole DSGE model
- Quasi-posterior uses only the IRF-matching information
- Strong identification: concentrates near $\hat{\theta}^{MD}$
- Weak identification: remains diffuse in weak directions

Backup: weak-GMM limit

Monte Carlo Design

- Weak-identification design adapted from Canova and Sala (2009)

In deviations from steady state, with i.i.d. Gaussian shocks,

$$\begin{aligned}
\text{IS:}\quad & y_t = \psi_1\,\mathbb{E}_t y_{t+1} + \psi_2\,(i_t - \mathbb{E}_t \pi_{t+1}) + e_{1t}, \\
\text{NKPC:}\quad & \pi_t = \psi_3\,\mathbb{E}_t \pi_{t+1} + \psi_4\, y_t + e_{2t}, \\
\text{Policy:}\quad & i_t = \psi_5\,\mathbb{E}_t \pi_{t+1} + e_{3t},
\end{aligned}$$

- Dual-mandate welfare loss (Rudebusch and Svensson, 1999)

$$W(a,\psi) = -\{\mathrm{Var}_\psi(\pi; a) + \mathrm{Var}_\psi(y; a) + 0.1\ \mathrm{Var}_\psi(\Delta i; a)\}$$

- Welfare-relevant parameters: ψ_2 and ψ_4
- Identification regimes: strong and weak
- Benchmark: regret relative to oracle policy

Monte Carlo: Weak Identification

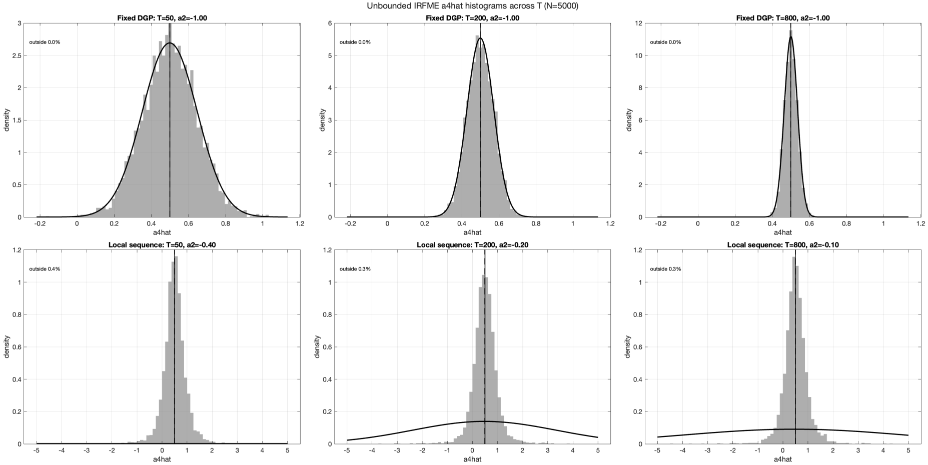

Sampling distributions of $\hat{\psi}_4$ under strong and weak identification

Monte Carlo Result

- Quasi-Bayes reduces average plug-in regret by 33.4%
- Quasi-Bayes leads to higher welfare in 67.5% of samples
- The two policies differ in approximately 81% of samples

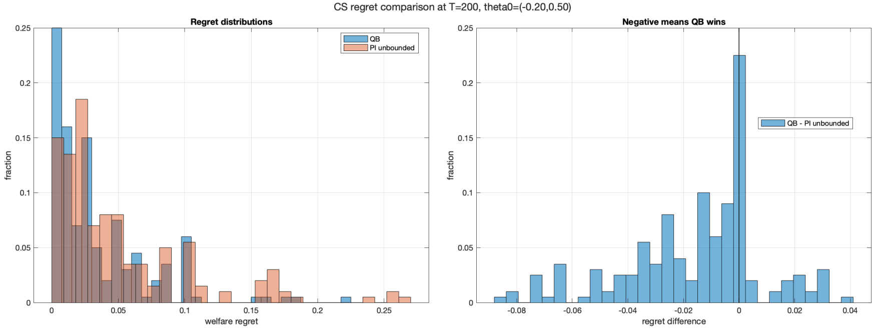

Regret relative to the oracle policy

Empirical Application: Medium-Scale NK Model

- Medium-scale New Keynesian model with standard nominal and real frictions (Christiano et al., 2005)
- Estimated by impulse-response matching to a monetary VAR
- Sample: US quarterly data, 1965:Q3–1995:Q3
- Relative IRF-matching information criterion of Hall et al. (2012) selects $H = 7$
- Policy: backward-looking Taylor policy

$$i_t = \rho_r i_{t-1} + (1 - \rho_r)\{\varphi_\pi \pi_{t-1} + \varphi_y y_{t-1}\}, \qquad a = (\varphi_\pi, \varphi_y, \rho_r)$$

- Welfare: dual-mandate

$$W(a, \psi) = -\{\mathrm{Var}(\pi) + \mathrm{Var}(y) + 0.1\ \mathrm{Var}(\Delta i)\}.$$

Medium-Scale NK Model: Quasi-Posterior

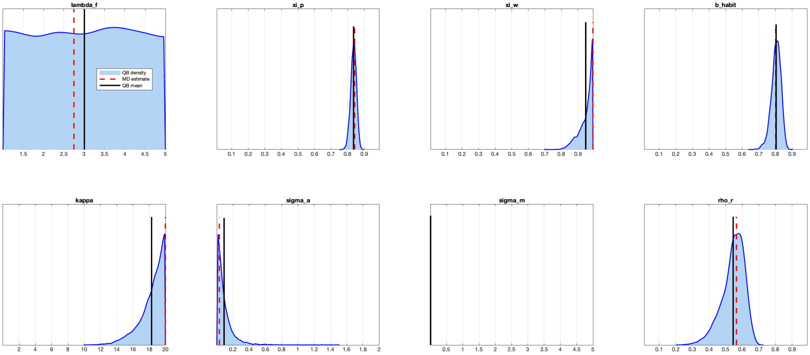

Shaded: quasi-posterior density; dashed: MD estimate; solid: QB mean

Medium-Scale NK Model: Quasi-Posterior

- Several parameters are well pinned down relative to the prior
  - price stickiness ξ_p
  - habit b
  - sample-policy smoothing ρ_r
- Other dimensions remain weakly disciplined or sit at economically meaningful bounds
  - price markup λ_f remains very diffuse
  - wage stickiness ξ_w hits its upper bound
  - investment adjustment cost κ hits its upper bound
  - monetary-shock scale σ_m hits its lower bound
- Their uncertainty is transmitted to policy through the welfare ranking over Taylor rules
  - price stickiness ξ_p and wage stickiness ξ_w change inflation and wage dynamics
  - monetary-shock scale σ_m changes the size of policy disturbances
  - price markup λ_f mainly shifts the markup block and has a limited policy effect here

Medium-Scale NK Model: Welfare Surface

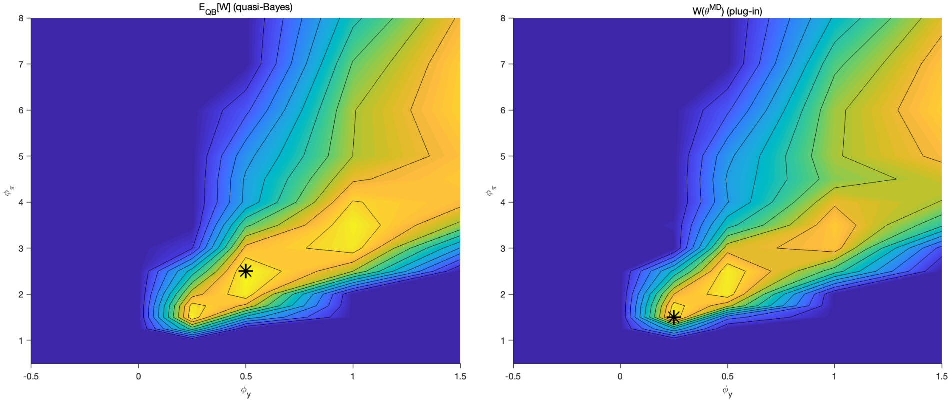

Left: QB posterior-mean welfare; right: plug-in welfare at $\hat{\theta}^{MD}$

Medium-Scale NK Model: Rule Disagreement

The selected policies $a = (\varphi_\pi, \varphi_y, \rho_r)$ are

$$\widehat{a}^{\mathrm{PI}} = (1.50,\ 0.25,\ 0.90)$$
$$\widehat{a}^{\mathrm{QB}} = (2.50,\ 0.50,\ 0.95)$$

- Plug-in reflects Taylor-style inflation response with high smoothing (Taylor, 1993; Clarida et al., 2000)
- Quasi-Bayes favors stronger stabilization
  - some draws imply more persistent inflation and output dynamics
  - stronger responses reduce inflation and output volatility

Backup: caveats

Conclusion

- Weak identification is a policy problem
  - weakly identified parameters can rank policies differently
- Plug-in policy evaluates welfare at one parameter value
  - sampling uncertainty is discarded before the policy step
- Quasi-Bayes carries structural uncertainty into welfare
  - choose the policy with lowest quasi-posterior expected loss
- Evidence
  - Monte Carlo: quasi-Bayes reduces regret
  - medium-scale NK model: selected policies disagree

References I

Andrews, I. and Mikusheva, A. (2022). Optimal decision rules for weak GMM. *Econometrica*, 90(6):2547–2594.

Canova, F. and Sala, L. (2009). Back to square one: Identification issues in DSGE models. *Journal of Monetary Economics*, 56(4):431–449.

Chari, V. V., Kehoe, P. J., and McGrattan, E. R. (2009). New Keynesian models: Not yet useful for policy analysis. *American Economic Journal: Macroeconomics*, 1(1):242–266.

Chernozhukov, V. and Hong, H. (2003). An MCMC approach to classical estimation. *Journal of Econometrics*, 115(2):293–346.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45.

Clarida, R., Galí, J., and Gertler, M. (2000). Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics*, 115(1):147–180.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton University Press, Princeton, 2 edition.

Giannoni, M. P. and Woodford, M. (2005). Optimal inflation-targeting rules. In Bernanke, B. S. and Woodford, M., editors, *The Inflation-Targeting Debate*, pages 93–162. University of Chicago Press, Chicago.

Giannoni, M. P. and Woodford, M. (2017). Optimal target criteria for stabilization policy. *Journal of Economic Theory*, 168:55–106.

Hall, A. R., Inoue, A., Nason, J. M., and Rossi, B. (2012). Information criteria for impulse response function matching estimation of DSGE models. *Journal of Econometrics*, 170(2):499–518.

References II

Hansen, L. P. and Sargent, T. J. (2008). *Robustness*. Princeton University Press, Princeton.

Mavroeidis, S., Plagborg-Møller, M., and Stock, J. H. (2014). Empirical evidence on inflation expectations in the New Keynesian phillips curve. *Journal of Economic Literature*, 52(1):124–188.

Rudebusch, G. D. and Svensson, L. E. O. (1999). Policy rules for inflation targeting. In Taylor, J. B., editor, *Monetary Policy Rules*, pages 203–262. University of Chicago Press, Chicago.

Schmitt-Grohé, S. and Uribe, M. (2007). Optimal simple and implementable monetary and fiscal rules. *Journal of Monetary Economics*, 54(6):1702–1725.

Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy*, 39:195–214.

Backup: Weak-GMM Limit Experiment

The normalized moment process satisfies

$$m_T(\cdot) \rightsquigarrow Z(\cdot), \qquad Z \sim P_\mu,$$

where Z is a Gaussian process with unknown mean function μ and consistently estimable covariance kernel.

The Gaussian limit induces a quasi-likelihood of the form

$$\mathcal{L}_Q(\theta; Z) \propto \exp\{-Q(\theta)/2\}.$$

Its sample analogue is

$$Q_T(\theta) = m_T(\theta)'\widehat{\Sigma}(\theta)^{-1} m_T(\theta),$$

which yields the quasi-posterior used in the main text.

Andrews and Mikusheva (2022) show that this quasi-Bayes construction is the diffuse-prior Bayes rule in the weak-GMM limit experiment.

Back

Backup: Caveats

- Medium-scale NK result is qualitative
  - empirical object is rule disagreement
  - not a large welfare-gain claim
- Weak directions can inherit prior weight
  - diffuse prior avoids adding tight curvature
  - prior sensitivity remains relevant under weak identification
- Quasi-Bayes credible sets are Bayesian objects
  - they need not have frequentist coverage

Back