A Decision-Theoretic Approach to Monetary Policy under Weak Identification

Jorge de la Cal Medina

University of Amsterdam, Tinbergen Institute

June 23, 2026

Abstract

DSGE models with different structural parameters can be observationally equivalent for the moments used in limited-information estimation. This is consequential for policy: parameter values that generate the same moments can imply different welfare rankings of policy rules. We study the choice of an interest-rate rule from a statistical decision-theoretic perspective. We apply a quasi-Bayes decision rule that minimizes posterior expected loss. The rule averages welfare over the structural values the moments leave plausible. We illustrate the construction in a Monte Carlo design and in a medium-scale New-Keynesian DSGE model estimated by impulse-response matching.

1 Introduction

DSGE models provide a structural environment in which alternative policies can be evaluated in order to decide upon an optimal policy. The model specifies how private decisions, frictions, shocks, and policy interact in equilibrium, so that a change in the policy can be mapped into outcomes and welfare. This interaction, however, depends on the underlying structural parameters. In empirical work these parameters can be disciplined by selected implications of the model, such as moments, auto-covariances, spectra, or impulse responses. A central difficulty is that this mapping need not be one-to-one: different structural parameter vectors can generate the same, or nearly the same, empirical objects.

This distinction is consequential for policy. Chari et al. (2009) emphasize that reduced-form objects can fit the data while admitting several structural interpretations with different policy implications. In the DSGE identification literature, Canova and Sala (2009) show that impulse-response matching can leave structural parameters unidentified because parameters may drop out of the solution, enter the matched responses only through combinations, or generate objective functions with little curvature. Thus parameter vectors that are hard to distinguish empirically may have different welfare properties, and a DSGE model can be disciplined by the moments used for estimation while leaving unresolved the structural state that matters for policy.

This paper studies that problem for the choice of an operational interest-rate rule in a New Keynesian environment. In such models, the welfare ranking of policy actions depends on the parameters governing nominal rigidities, real frictions, and shock propagation (Woodford, 2003; Galí, 2015). If those parameters are weakly identified by the moments used in estimation, plug-in policy choice treats as known precisely the object the data have failed to pin down. The relevant question is then not only which parameter vector best fits the data, but which policy performs well across the structural values the data leave plausible.

The argument of this paper is that, once parameter uncertainty of this kind is taken seriously, the selection of a policy is not an estimation problem followed by an optimization but a statistical decision problem in the sense of Wald (1950). The action is an operational rule, the state is the structural parameter that generated the data, and the loss is the welfare cost incurred by the rule at that state. The object of interest is not the parameter value that best fits the data but the rule that performs well across the parameter values the data leave plausible. In the local experiments associated with weak identification, the directions left undetermined by the moments cannot be treated as consistently estimable. Policy choice then requires an explicit criterion for trading off expected loss across states. We use the quasi-Bayes construction of Andrews and Mikusheva (2022) to form this trade-off; the resulting operational policy maximizes quasi-posterior expected welfare.

Relation to the literature on robust policy. This paper is closely related to robust monetary policy literature. Here, "robustness" denotes several distinct programs. In the Hansen–Sargent tradition (Hansen and Sargent, 2008), the policymaker knows the structural parameters but distrusts the probability law governing outcomes; specification doubts are disciplined by a set of models surrounding a benchmark, and the rule is evaluated under the least favorable member of that set, yielding a max-min criterion.

Giannoni and Woodford (2005, 2017) pursue a different notion: within a correctly specified equilibrium they derive robust targeting criteria that implement the optimal allocation without detailed knowledge of the disturbance processes, and a related literature studies robustness to departures from rational expectations (Woodford, 2010).

This paper addresses a third and complementary source of doubt: weak empirical discipline on the structural parameters that matter for policy. The issue is not that the policymaker distrusts the model's probability law or seeks a rule that is robust to unspecified disturbances. It is that the empirical moments used to discipline the DSGE model may leave several welfare-relevant structural states plausible in the weak-GMM limit.

Contributions. The paper makes three contributions. First, it shows that weak identification in limited-information DSGE estimation has direct policy consequences: structural parameters that are observationally equivalent for the matched moments can imply different welfare rankings of operational interest-rate rules. Plug-in policy choice therefore inherits the instability of the point estimate. Second, it formulates the choice of a policy rule as a statistical decision problem. The action is the rule, the state is the structural parameter, and the loss is the welfare cost of using the rule at that state. This formulation makes clear that plug-in is a degenerate Bayes rule and that, under weak identification, policy choice requires an explicit ordering of risk across the structural values left plausible by the data. Third, following Andrews and Mikusheva (2022), it constructs a feasible quasi-Bayes rule to choose a policy. The rule minimizes posterior expected loss and coincides with plug-in under strong identification, but under weak identification it averages welfare over the relevant set of plausible structural values. We illustrate the mechanism in a Monte-Carlo design, where the welfare gain can be measured against the truth, and in a medium-scale DSGE model estimated by impulse-response matching, where the quasi-Bayes rule reveals how weak identification changes the policy recommendation.

Outline. Section 2 states the policy problem under a known model and isolates the dependence of the optimal rule on the structural parameters. Section 3 sets up estimation by the continuously-updating GMM criterion, distinguishes strong identification, weak identification, and misspecification, and derives the Gaussian limit. Section 4 formulates policy choice as a statistical decision problem on the limiting problem. Section 5 constructs the quasi-Bayes operational policy and establishes its optimality and its strong-identification limit. Section 6 reports the application.

2 The Policy Problem under Parameter Uncertainty

The economic model is indexed by a structural parameter $\psi \in \Psi$ and a policy action $a \in \mathcal{A}$. The action indexes a policy; in the monetary-policy application, $\mathcal{A}$ is a parametric family of interest-rate rules, such as Taylor rules $a = (\varphi_\pi, \varphi_y, \rho_r)$ that respond to inflation and the output gap with interest-rate smoothing. Each pair (ψ, a) induces an equilibrium, which the policymaker evaluates through the welfare function $W(a, \psi)$.

Assumption 1 (Action class and welfare)**.** *$\mathcal{A}$ is compact. For each $\psi \in \Psi$ the determinacy set $\mathcal{A}_\psi = \{a \in \mathcal{A} : (a, \psi)$ induces a unique stable equilibrium with finite welfare$\}$ is nonempty, and $W(\cdot, \psi)$ is continuous on $\mathcal{A}_\psi$. $W(a, \psi)$ is bounded on $\{(a, \psi) : a \in \mathcal{A}_\psi\}$.*

We focus on cases in which the parameters (ψ, a) determine the equilibrium objects entering $W(\cdot, \cdot)$, not the welfare function itself.

If ψ were known, the optimal operational policy would solve

$$a^\star(\psi) \in \arg\max_{a\in\mathcal{A}_\psi} W(a,\psi). \tag{1}$$

This is the conventional optimal-policy problem, whether the maximization is unrestricted, as in Ramsey analysis, or confined to an operational class of implementable rules (Woodford, 2003; Giannoni and Woodford, 2005, 2017; Schmitt-Grohé and Uribe, 2007). Even with ψ known, the problem is substantive because evaluating $W(a,\psi)$ requires solving the equilibrium induced by each candidate action.

The difficulty studied here is that ψ is not known. Different structural parameter values can imply different welfare rankings over the same policy actions. A rule that is optimal at one value of ψ need not be optimal at another. Policy choice under parameter uncertainty therefore requires a criterion for comparing rules across the structural values left plausible by the data.

Usually ψ is inferred from data. Let $X_T = \{x_t\}_{t=1}^T$ be the observed sample, with support $\mathcal{X}$ and law P. A decision rule $\delta : \mathcal{X} \to \mathcal{A}$ maps the sample into the action space. With loss $L(a,\psi) = -W(a,\psi)$, the rule is evaluated at θ by its risk

$$\int_{\mathcal{X}} L(\delta(x),\psi)\,dP(x). \tag{2}$$

The goal is to choose a rule with a desirable risk function. Which risk functions are attainable, and how they should be ordered, depends on how the data identify the structural parameters.

3 Identification and Inference

Let sample X_T follow law P_θ in a model $\mathcal{P} = \{P_\theta : \theta \in \Theta\}$. The maintained parameter is $\theta = (\psi,\alpha)$. The component ψ collects the policy-invariant structural parameters, while α describes the policy prevailing in the sample and therefore affects the law of observed data and moments. The counterfactual policy experiment holds ψ fixed and replaces α by the chosen action a; accordingly loss depends on θ only through ψ. The moments may carry information about both blocks, so identification is posed for the full vector θ.

Assumption 2 (Moment condition)**.** *There is a known function $\phi : \mathcal{X}\times\Theta \to \mathbb{R}^k$ such that, under the true law P_{θ_0}, the true value $\theta_0 = (\psi_0,\alpha_0)$ satisfies*

$$\mathbb{E}_{\theta_0}[\phi(x_t,\theta_0)] = \int_{\mathcal{X}} \phi(x_t,\theta_0)\,dP_{\theta_0}(x_t) = 0.$$

Inference proceeds through the sample-moment process,

$$m_T(\theta) = \frac{1}{\sqrt{T}} \sum_{t=1}^{T} \phi(x_t, \theta), \tag{3}$$

with which we can build the continuously-updating GMM criterion (Hansen, 1982; Hansen et al., 1996)

$$Q_T(\theta) = m_T(\theta)' \, \widehat{\Sigma}(\theta)^{-1} \, m_T(\theta) \tag{4}$$

with $\widehat{\Sigma}(\theta)$ a consistent estimator of $\lim_{T\to\infty} \mathrm{Var}_{\theta_0}(m_T(\theta))$. This gives the continuously-updating estimator (CUE), $\hat{\theta} = \arg\min_{\theta\in\Theta} Q_T(\theta)$.

Example. Impulse-response matching is the running example. It is a common limited-information way to estimate structural DSGE models without specifying the full likelihood (Christiano et al., 2005; Hall et al., 2012). Let $\hat{\gamma}_T \in \mathbb{R}^k$ collect the selected empirical impulse responses estimated from a reduced-form vector autoregression. If r responses are matched over H horizons, then $k = rH$. We assume $\hat{\gamma}_T$ is consistent for the population response vector $\gamma_0 \in \mathbb{R}^k$, and let $g(\theta) \in \mathbb{R}^k$ collect the corresponding model-implied responses. Under correct specification, the population responses satisfy $\gamma_0 = g(\theta_0)$. Estimation proceeds by minimising

$$m_T^{IRF}(\theta) = \sqrt{T}(\hat{\gamma}_T - g(\theta)),$$

weighted with $\widehat{V}$, a consistent estimator for the asymptotic covariance of $\sqrt{T}(\hat{\gamma}_T - \gamma_0)$. The procedure matches selected reduced-form dynamics rather than the full likelihood. Because $\widehat{V}$ does not depend on θ, the criterion is a fixed-weight minimum-distance criterion.

We discuss different identification cases. Consider the identified set

$$\Theta_0 = \big\{\theta \in \Theta : \mathbb{E}_{\theta_0}[\phi(x_t, \theta)] = 0\big\}, \tag{5}$$

the set of parameter values consistent with the population moment condition (Assumption 2) under the true law P_{θ_0}. We distinguish between strong identification, weak identification, and misspecification.

Strong identification. $\Theta_0 = \{\theta_0\}$ is a singleton: θ_0 is the unique value satisfying the population moment condition. Then ψ_0 is recovered from the population moments, and the population optimal policy is $a^\star(\psi_0)$.

Weak identification. The moment condition is close to zero, relative to sampling uncertainty, over a nontrivial part of the parameter space. Following Andrews and Mikusheva (2022, Section 2.1), we model this by sequences of data-generating processes that are local to an identification failure. At the limiting law, the moment condition holds at more than one value, so Θ_0 contains several parameter values. Let $\Psi_0 = \{\psi : (\psi, \alpha) \in \Theta_0\}$ be the corresponding set of structural components relevant for welfare. This component is therefore identified only up to the set Ψ_0, and the associated optimal policies form the set $\{a^\star(\psi) : \psi \in \Psi_0\}$.

Misspecification. No $\theta \in \Theta$ satisfies

$$\mathbb{E}_{P_0}[\phi(x_t, \theta)] = 0,$$

for a true limiting law P_0, so the identified set Θ_0, now taken under P_0, is empty. The relevant object is then a pseudo-true value minimizing the population criterion, which can depend on the weighting matrix (Hall and Inoue, 2003).

The remainder of the paper focuses on strong and weak identification under correct specification. Extending the policy problem to misspecification would require a separate pseudo-true or enlarged-state decision problem. Although relevant, that case does not fit the current framework and is left for future work.

Example (cont.). In impulse-response matching, strong identification corresponds to a unique θ_0 such that $\gamma_0 = g(\theta_0)$. Weak identification corresponds to multiple, or nearly multiple, structural values matching the same population responses, so that $\gamma_0 = g(\theta^*)$ for multiple $\theta^* \in \Theta_0$. Canova and Sala (2009) show that this can originate at several layers of the DSGE-to-criterion map: parameters may have little effect on the solution coefficients, the selected shocks or variables may carry little information about particular mechanisms, and the distance criterion may display ridges or low curvature. In rich models, different frictions can therefore generate similar conditional dynamics, so many structural values fit the reduced-form responses equally well while implying different ψ and hence different welfare rankings.

Inference on θ is conducted through the moment conditions. Under weak identification, regular normal approximations to GMM estimators can fail, but the scaled sample-moment process still admits a Gaussian-process limit (Andrews and Mikusheva, 2022, Sections 2.1–2.2). This limit experiment is the object through which we compare the large-sample risk of decision rules. Under mild conditions, the sample-moment process has a Gaussian limit regardless of whether the structural parameter is strongly or weakly identified:

$$m_T(\cdot) \rightsquigarrow Z(\cdot), \qquad Z \sim P_\mu, \tag{6}$$

where P_μ is the law of a Gaussian process with unknown mean function $\mu(\cdot)$ and covariance kernel $\Sigma(\cdot, \cdot)$. Following Andrews and Mikusheva (2022, Section 2.2), the covariance kernel is treated as known in the limit experiment because it is consistently estimable.

Example (cont.). For the fixed-weight IRF criterion, the specialization of (6) is a Gaussian shift process. The empirical IRF fluctuation contributes the Gaussian term with covariance V, while local population discrepancies between γ_0 and $g(\theta)$ contribute the mean function. Thus the limiting process has the form

$$Z^{IRF}(\theta) = \mu^{IRF}(\theta) + G, \qquad G \sim N(0, V),$$

with covariance kernel $\Sigma^{IRF}(\theta_1, \theta_2) = V$.

We use the asymptotic representation theorem of Andrews and Mikusheva (2022, Theorem 1) to work with this limit experiment. The theorem shows that the risk functions attainable in the limit experiment lower-bound attainable asymptotic risk functions in the original GMM experiment. Hence asymptotic optimality can be studied by solving the decision problem in the limit experiment.

A decision rule in the limit experiment is a measurable map $\delta_G : \mathcal{Z} \to \mathcal{A}$, where $\mathcal{Z}$ denotes the sample space of the Gaussian process Z. At a state $\theta^* = (\psi^*, \alpha^*)$ and mean function μ satisfying $\mu(\theta^*) = 0$, it is evaluated by

$$\mathbb{E}_{P_\mu}[L(\delta_G(Z), \psi^*)] \,. \tag{7}$$

This is the large-sample analogue of the finite-sample risk in (2): the expectation is taken over the limiting distribution of the sample-moment process, while the loss is evaluated at the structural component that enters welfare.

Under strong identification, the limiting zero set contains a single structural value. The risk in (7) is then evaluated at that point, and a rule that converges to the population optimum $a^\star(\psi_0)$ attains the limit risk. This is the usual plug-in logic.

Under weak identification, the limiting zero set contains several structural values. The object in (7) is then a risk function over those values: a rule that minimizes risk at one value of ψ need not minimize it at another. There is therefore no uniformly best decision rule in general. A decision criterion must specify how risks are ordered across the structural values left in the limiting experiment; the next section supplies this ordering.

4 The Decision Problem in the Gaussian Limit

Choosing an operational rule under weak identification is therefore a problem of ranking the Gaussian risk functions defined in (7). Since no rule is uniformly best in general, a criterion is needed to order risk across the structural values left in the limiting experiment.

A rule δ_G *dominates* δ'_G if its Gaussian risk is no larger at every $\theta \in \Theta_0$, and strictly smaller for some θ, and δ_G is *admissible* if no rule dominates it. An inadmissible rule is uniformly weakly worse than some alternative and can be discarded without any stand on how risk is weighted across states. Under weak identification, however, admissibility can leave a large class of rules with different risk profiles over Θ_0, so a complete ordering is needed.

To make this trade-off explicit, we follow Andrews and Mikusheva (2022) and adopt a Bayes criterion. Let π be a prior on Θ_0, which weights risks across the structural values in the limiting experiment. For each realization z of the Gaussian process, the Bayes action minimizes posterior expected loss,

$$\delta_G^{\star}(z) \in \arg\min_{a \in \mathcal{A}} \int_{\Theta_0} L(a, \psi)\, d\pi(\theta \mid z). \tag{8}$$

Bayes rules are useful here because they provide an admissible way to order risk functions. Under the usual regularity conditions for Bayes decision problems, Bayes rules with full-support priors are admissible in the limit experiment: they are not uniformly dominated in Gaussian risk. Complete-class results can deliver converse statements under additional conditions.

The plug-in rule is the degenerate case of this criterion. If the prior places point mass on a single value θ_0, posterior expected loss reduces to $L(a, \psi_0)$, and the rule chooses the population optimum $a^{\star}(\psi_0)$.

Criteria that avoid prior weights need not resolve this choice. Minimax and minimax-regret reduce each risk function to its worst value over Θ_0. When several rules share the same worst-case value, the criterion is silent among them. Moreover, a rule that minimizes worst-case risk need not be admissible, so a minimax criterion may select a dominated rule. The next section implements this Bayes criterion using the quasi-Bayes construction of Andrews and Mikusheva (2022).

5 Quasi-Bayes Policy Choice

To implement the Bayes criterion, we need the posterior distribution over structural values obtained by combining a prior with the observed moment process. In the Gaussian limit, a candidate value θ is evaluated under the restriction that the mean function vanishes at that value, $\mu(\theta) = 0$. The unrestricted part of the mean function is an infinite-dimensional nuisance object. In the Gaussian limit experiment, the normalized moment process has a consistently estimable covariance kernel. We place independent priors on the structural parameter and the nuisance mean function, and use a proportional Gaussian-process prior for the nuisance component. The diffuse-prior limit then yields a marginal likelihood contribution for θ, up to terms absorbed into the prior, given by the quasi-likelihood

$$\mathcal{L}_Q(\theta; Z) \propto \exp\{-Q(\theta)/2\}, \tag{9}$$

where $Q(\theta) = Z(\theta)'\Sigma(\theta, \theta)^{-1} Z(\theta)$; determinant terms that do not depend on the realized process are absorbed into the prior (Andrews and Mikusheva, 2022, equation (8)). This is a quasi-likelihood because it treats the Gaussian approximation to the moment process as the basis for updating θ, rather than specifying a likelihood for the original sample X_T.

The feasible quasi-Bayes procedure replaces the limiting objects Z and Σ by the normalized sample moments m_T and the estimated covariance $\widehat{\Sigma}$:

$$Q_T(\theta) = m_T(\theta)'\widehat{\Sigma}(\theta)^{-1}m_T(\theta), \tag{10}$$

and the quasi-posterior

$$\pi_{\mathrm{QB}}(\theta \mid X_T) = \frac{\exp\{-Q_T(\theta)/2\}\,\pi(\theta)}{\int_\Theta \exp\{-Q_T(u)/2\}\,\pi(u)\,du}. \tag{11}$$

This object is a quasi-posterior, not a Bayesian posterior, unless the GMM criterion is itself a correctly specified likelihood (Chernozhukov and Hong, 2003). Under weak identification its dependence on the prior is substantive: the criterion concentrates the quasi-posterior along strongly identified directions, while weakly identified directions inherit prior mass. We therefore use a diffuse prior to avoid adding tight prior curvature beyond the information in the moments.

The quasi-Bayes policy chooses the operational rule that minimizes posterior expected loss:

$$\delta_{\mathrm{QB}}(X_T) \in \arg\min_{a\in\mathcal{A}} \int_\Theta L(a,\psi)\,\pi_{\mathrm{QB}}(\theta \mid X_T)\,d\theta, \qquad \theta = (\psi,\alpha). \tag{12}$$

We write $\widehat{a}_{\mathrm{QB}} = \delta_{\mathrm{QB}}(X_T)$ for the realized action. For each draw $\theta^{(m)} = (\psi^{(m)}, \alpha^{(m)})$ from $\pi_{\mathrm{QB}}(\cdot \mid X_T)$, welfare is evaluated by holding $\psi^{(m)}$ fixed and replacing the sample policy component $\alpha^{(m)}$ with the candidate rule a.

The simple plug-in rule is $\widehat{a}_{\mathrm{PI}} = a^\star(\widehat{\psi})$, where $\hat{\theta} = (\widehat{\psi}, \widehat{\alpha})$. It is equivalent to replacing (11) by a point mass at $\hat{\theta} = (\widehat{\psi}, \widehat{\alpha})$ and then evaluating welfare through $\widehat{\psi}$. It therefore discards the weakly identified directions. Under strong identification the quasi-posterior concentrates at ψ_0, so $\widehat{a}_{\mathrm{QB}}$ and $\widehat{a}_{\mathrm{PI}}$ have the same asymptotic limit (Chernozhukov and Hong, 2003); under weak identification (12) instead averages welfare over structural values that the moments leave plausible.

The decision rule (12) is the feasible quasi-posterior version of the posterior-loss rule (8). It uses the full parameter space and the sample criterion Q_T, rather than the unknown identified set and the limiting Gaussian process. Andrews and Mikusheva (2022, Theorem 5 and Corollary 1) show that this feasible construction has the same asymptotic decision content as the corresponding infeasible quasi-posterior rule in the Gaussian weak-GMM limit. Hence the rule chosen from (12) is not an ad hoc averaging device: it is the sample analogue of the Bayes rule in the weak-GMM limit.

Example (cont.). For impulse-response matching,

$$\pi_{\mathrm{QB}}(\theta \mid X_T) \propto \exp\left[-\frac{T}{2}\{\hat{\gamma}_T - g(\theta)\}'\widehat{V}^{-1}\{\hat{\gamma}_T - g(\theta)\}\right]\pi(\theta). \tag{13}$$

Given draws from (13), the operational rule is computed by maximizing average welfare over $a \in \mathcal{A}$, by grid search or numerical optimization.

6 Empirical Application

We illustrate the quasi-Bayes operational policy in two complementary settings that share the same machinery: a structural model is estimated by matching impulse responses, the quasi-posterior (11) is sampled, and the operational rule (12) is computed on a grid and compared with the plug-in rule. The two settings answer different questions. The first is a controlled Monte-Carlo laboratory in the weakly identified design of Canova and Sala (2009): because the truth is known, we can measure the welfare *gain* of the quasi-Bayes rule, but the calibration is not meant as a literal policy prescription. The second is a medium-scale New Keynesian model following Christiano et al. (2005) in its nominal and real frictions, estimated on US data, where the truth is unknown and the question is *qualitative*: how the quasi-Bayes rule departs from the plug-in rule when identification is weak. Read together they make the controlled-to-real case: the mechanism is established where it can be measured, then exhibited where it matters.

6.1 Monte-Carlo simulation

The design of Canova and Sala (2009) is a small New Keynesian model in which the structural slopes are weakly identified from impulse responses. Under i.i.d. shocks the model reduces to a static system in the IS and NKPC slopes $\psi = (\psi_2, \psi_4)$; the impulse responses to the policy shock identify ψ through a Jacobian whose determinant is proportional to the IS slope ψ_2, so identification degrades as $\psi_2 \to 0$, at which point the NKPC slope ψ_4 becomes unidentified. The data-generating values $\psi_0 = (\psi_2, \psi_4) = (-0.2, 0.5)$ lie in this weakly identified region. Appendix A gives the model and its solution. The policymaker chooses a contemporaneous Taylor rule $a = (\varphi_\pi, \varphi_y)$ under the dual-mandate loss

$$W(a, \psi) = -\big[\mathrm{Var}_\psi(\pi; a) + \lambda_x \mathrm{Var}_\psi(y; a) + \lambda_R \mathrm{Var}_\psi(\Delta i; a)\big], \tag{14}$$

with $\lambda_x = 1$ and $\lambda_R = 0.1$, using the Rudebusch–Svensson quadratic dual-mandate loss form with interest-rate smoothing (Rudebusch and Svensson, 1999). The last term penalizes instrument volatility. It prevents the controlled comparison from rewarding rules that stabilize inflation and output only by generating sharp movements in the nominal rate.

At the truth ψ_0 the welfare-optimal rule is interior, $a^\star(\psi_0) = (0.50, 1.25)$, so the comparison is not contaminated by a boundary; imposing a Taylor-principle floor $\varphi_\pi \geq 1$

would exclude the truth-optimal rule and is therefore not imposed.[1]

We draw $N = 200$ samples of length $T = 200$, and for each sample compute the quasi-Bayes and plug-in rules and their regret relative to the oracle $a^{\star}(\psi_0)$. Table 1 summarizes the result: the quasi-Bayes rule eliminates one third of the plug-in's expected regret and improves on it in two thirds of samples. The gain is not an artifact of the dual-mandate weight; a sweep over λ_x leaves it above one half even at the Woodford micro-founded value $\lambda_x = 0.1$. It behaves as the theory predicts across sample sizes: under fixed-parameter asymptotics it decays as the quasi-posterior concentrates (the strong-identification limit of Section 5), whereas along a weak-identification sequence $\psi_2 = -c/\sqrt{T}$ it is flat and does not vanish (Figure 1). The finite-sample distortion in the IRF estimate is common to both procedures, so the comparison isolates the decision rule rather than the consistency of the point estimator.

Table 1: Canova–Sala Monte Carlo ($N = 200$, $T = 200$), regret relative to the oracle rule $a^{\star}(\psi_0) = (0.50, 1.25)$.

Quasi-Bayes share of plug-in regret eliminated	33.4%
Samples in which $\widehat{a}_{\mathrm{QB}}$ beats $\widehat{a}_{\mathrm{PI}}$	67.5%
Samples in which the two rules disagree	$\approx 81\%$

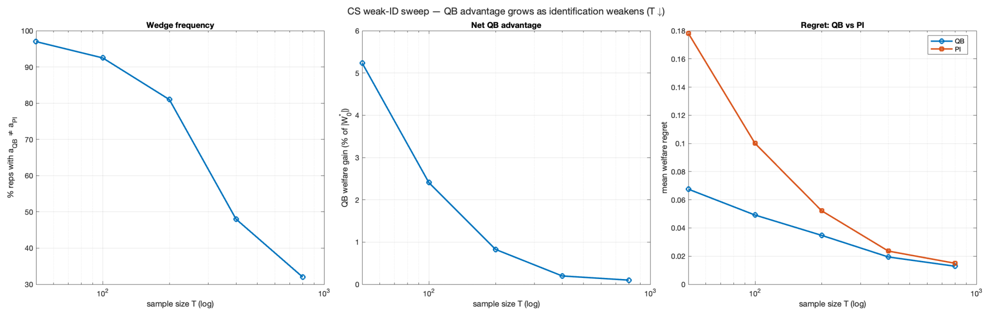

Figure 1: Quasi-Bayes regret reduction across sample sizes. Under fixed-parameter asymptotics the advantage decays toward the strong-identification limit; along the weak-identification (Pitman-drift) sequence it remains bounded away from zero.

6.2 A medium-scale New Keynesian model on US data

We estimate a medium-scale New Keynesian model following Christiano et al. (2005) in its central nominal and real frictions: Calvo price and wage setting with indexation, habit formation, investment adjustment costs, variable capital utilisation, and a working-capital cost channel. The empirical implementation is not a replication of Christiano

[1] The design is static, so the determinacy denominator is positive for every admissible rule (Appendix A) and the floor has no equilibrium justification here. Imposing it nonetheless reverses the ranking for $T \geq 200$, which we report as a quantified caveat rather than suppress.

et al. (2005): it is estimated by impulse-response matching to a recursively identified monetary VAR on a reduced set of US observables over 1965:Q3–1995:Q3. The relative IRF-matching information criterion of Hall et al. (2012) selects the matching horizon $H = 7$. Table 2 reports the minimum-distance estimate $\hat{\theta}$ alongside the quasi-posterior mean and standard deviation. The table and the marginal densities in Figure 2 show why this is a useful real-data illustration of the decision problem. Some directions are tightly disciplined by the monetary-shock responses: the price Calvo parameter, habit, and the historical-rule smoothing coefficient have concentrated quasi-posterior marginals. Other directions are not. The price mark-up is nearly flat over the maintained prior range, the wage Calvo parameter and investment adjustment cost are pushed to their upper bounds by the minimum-distance criterion, and the monetary-shock scale is at its lower bound. The point estimate is therefore a poor summary of the structural uncertainty relevant for policy. The quasi-Bayes calculation uses the full quasi-posterior distribution over these structural values rather than selecting a rule at the single minimum-distance point. Appendix B details the data, the VAR identification, the matching objective, horizon selection, the sampler, and the policy grid.

Table 2: IRFME estimation of the medium-scale New Keynesian model, $H = 7$.

Parameter	MD $\hat{\theta}$	QB mean	QB std
λ_f (price mark-up)	2.75	3.00	1.15
ξ_p (Calvo price)	0.84	0.84	0.02
ξ_w (Calvo wage)	$0.99^{\dagger}$	0.95	0.04
b (habit)	0.80	0.80	0.03
κ (investment adj. cost)	$20.0^{\dagger}$	18.3	1.50
σ_a (utilisation)	0.03	0.09	0.10
σ_m (monetary shock s.d.)	$0.01^{\ddagger}$	0.01	0.002
ρ_r (smoothing)	0.57	0.55	0.07

† at upper bound; ‡ at lower bound.

The policymaker chooses a backward-looking Taylor rule $a = (\varphi_\pi, \varphi_y, \rho_r)$ under the same dual-mandate loss used in Section 6.1, equation (14), with $\lambda_x = 1$ and $\lambda_R = 0.1$. This is a deliberately simple operational class, close to the Taylor-rule tradition: a systematic response to inflation and real activity, with possible interest-rate smoothing. Taylor's original rule corresponds to a long-run inflation response around 1.5 and an output response around 0.5 (Taylor, 1993), while empirical reaction-function estimates commonly include substantial interest-rate smoothing and distinguish pre- and post-Volcker inflation responses (Clarida et al., 2000). The plug-in and quasi-Bayes rules selected on our grid are

$$\widehat{a}_{\mathrm{PI}} = (1.50,\ 0.25,\ 0.90), \qquad \widehat{a}_{\mathrm{QB}} = (2.50,\ 0.50,\ 0.95). \tag{15}$$

The plug-in rule is close to a conventional inertial Taylor rule. Its inflation coefficient is exactly the Taylor benchmark, its output coefficient is positive but lower than the original Taylor value, and its smoothing coefficient is high. The quasi-Bayes rule is not obtained by estimating different policy-rule coefficients from the data. The coefficients in a are controls. What changes is the criterion used to choose the control: plug-in maximizes welfare at $\hat{\theta}$, while quasi-Bayes maximizes welfare averaged over the quasi-posterior distribution of the structural parameters. Averaging over that distribution selects a rule with a stronger inflation response, a Taylor-sized output response, and slightly more inertia.

This rule disagreement is the empirical object. The structural uncertainty in Figure 2 is concentrated in economically relevant frictions and scales rather than in the policy coefficients themselves. When those structural draws are carried through the model solution and the welfare calculation, the ranking over candidate Taylor rules changes. Figure 3 shows the resulting policy surfaces. The plug-in surface is the welfare ranking at $\hat{\theta}$; the quasi-Bayes surface is the quasi-posterior mean welfare ranking. Their optima lie at different locations in the rule space. The CEE application therefore illustrates the paper's decision problem in a medium-scale monetary model: the data do not identify all welfare-relevant structural dimensions sharply, and the way one integrates over that uncertainty affects the operational rule chosen from a familiar monetary policy class.

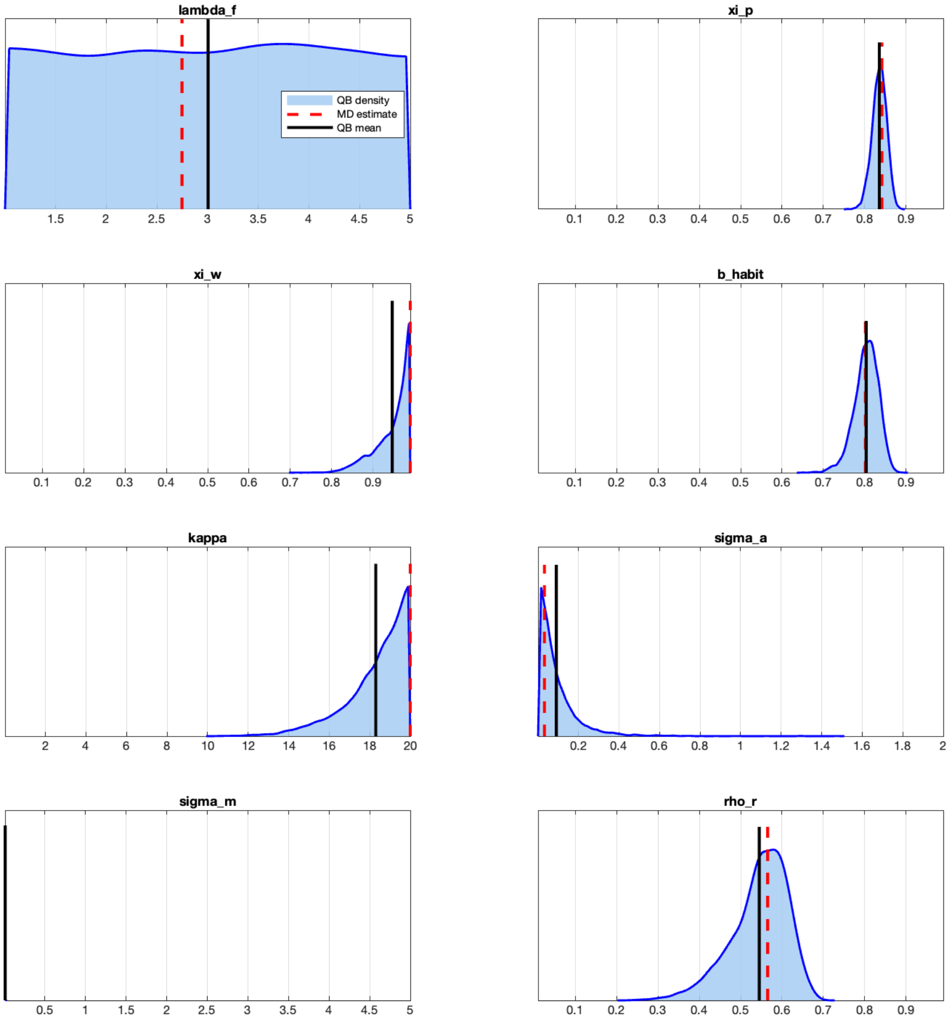

Figure 2: Quasi-posterior marginal densities of the eight structural parameters (shaded), with the minimum-distance estimate (dashed) and the quasi-posterior mean (solid). The impulse responses discipline some directions, such as ξ_p, b and the historical-rule smoothing coefficient ρ_r, but leave other directions weakly informed: λ_f is close to flat over the prior range, ξ_w and κ are pushed to upper bounds by the point criterion, and σ_m sits at its lower bound.

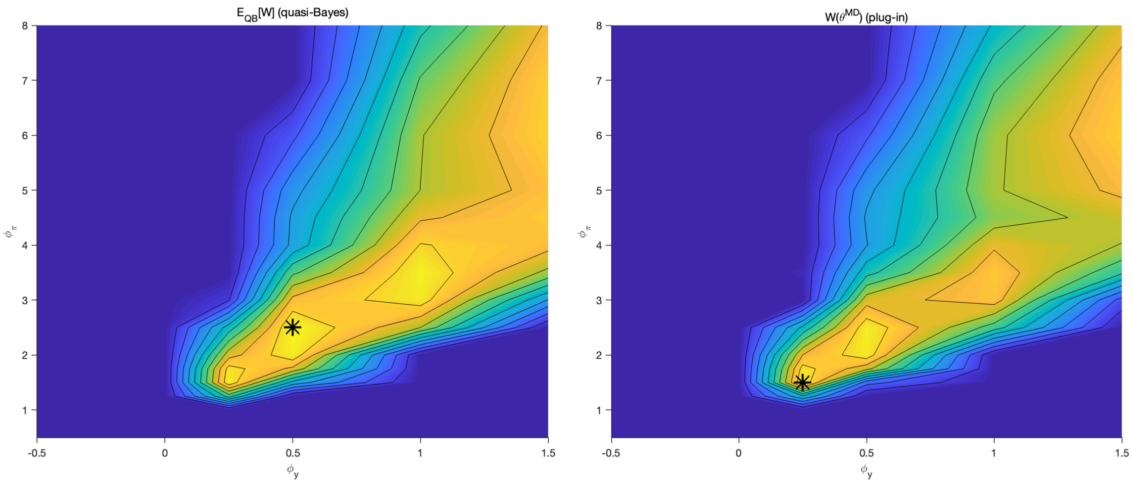

Figure 3: Welfare over the policy grid, profiled over the smoothing coefficient ρ_r. Left: quasi-Bayes posterior-mean welfare $E_{QB}[W]$; right: plug-in welfare $W(\hat{\theta})$. Markers: $\widehat{a}_{\mathrm{QB}}$ (star), $\widehat{a}_{\mathrm{PI}}$ (square). The two criteria rank the same Taylor-rule grid differently and select different operational rules.

A The Monte-Carlo model

This appendix records the model used in Section 6.1, its solution under the data-generating process, and its solution under the contemporaneous Taylor rule imposed in the policy experiment.

A.1 Model and static reduction

The design follows Canova and Sala (2009). In deviations from steady state,

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

with shocks $e_t = (e_{1t}, e_{2t}, e_{3t})' \sim \text{i.i.d.}\,\mathcal{N}(0, \Sigma_e)$, $\Sigma_e = \operatorname{diag}(\sigma_1^2, \sigma_2^2, \sigma_3^2)$. The structural block relevant for welfare is the pair of slopes $\psi = (\psi_2, \psi_4)$ (the interest-rate semi-elasticity of demand and the Phillips-curve slope); (ψ_1, ψ_3, ψ_5) are forward-looking weights and ψ_5 is the policy response to expected inflation.

Because the shocks are i.i.d., all conditional expectations of future variables are zero, $\mathbb{E}_t y_{t+1} = \mathbb{E}_t\pi_{t+1} = 0$, and the system collapses to the static form

$$\begin{aligned} y_t &= \psi_2\, i_t + e_{1t}, \\ \pi_t &= \psi_4\, y_t + e_{2t}, \\ i_t &= e_{3t}. \end{aligned}$$

The reduction annihilates the forward-looking weights (ψ_1, ψ_3, ψ_5): they drop out of the solution, mirroring the Canova–Sala observation that such coefficients are not identified from the model's dynamics. In particular the systematic policy response $\psi_5\,\mathbb{E}_t\pi_{t+1}$ vanishes, leaving $i_t = e_{3t}$.

A.2 Solution under the DGP

With $i_t = e_{3t}$, static reduction gives $Y_t = (y_t, \pi_t, i_t)' = c_\theta + A_0 e_t$ with

$$A_0 = \begin{pmatrix} 1 & 0 & \psi_2 \\ \psi_4 & 1 & \psi_2\psi_4 \\ 0 & 0 & 1 \end{pmatrix}. \tag{16}$$

This is the data-generating process (true values $\psi_2 = -0.2$, $\psi_4 = 0.5$). Estimation matches the impulse responses to the policy shock e_{3t} — the third column of A_0 — using its output and inflation entries,

$$m(\psi) = \sigma_3\,(\psi_2,\ \psi_2\psi_4)'. \tag{17}$$

The Jacobian of the matched moments is

$$\frac{\partial m}{\partial \psi} = \sigma_3 \begin{pmatrix} 1 & 0 \\ \psi_4 & \psi_2 \end{pmatrix}, \qquad \det\left(\frac{\partial m}{\partial \psi}\right) = \sigma_3^2\,\psi_2. \tag{18}$$

Identification strength is therefore governed by the IS slope ψ_2: as $\psi_2 \to 0$ the Jacobian becomes singular and ψ_4 is unidentified — a policy shock that does not reach output ($\psi_2 = 0$) leaves the responses uninformative about the Phillips-curve slope. The truth $\psi_2 = -0.2$ sits in this weakly identified region. Holding ψ_2 fixed, the parameter is point-identified and standard asymptotics apply (the quasi-Bayes advantage decays in T); the Pitman-drift sequence $\psi_2 = -c/\sqrt{T}$ instead keeps the design on the weak-identification frontier, where the advantage does not vanish.

A.3 Solution under the imposed Taylor rule

A forward-looking rule is inert under the static reduction, so to study operational policy we replace the policy equation by a contemporaneous Taylor rule

$$i_t = \varphi_\pi \pi_t + \varphi_y y_t + e_{3t}. \tag{19}$$

Substituting the NKPC into (19) gives $i_t = \kappa\, y_t + \varphi_\pi e_{2t} + e_{3t}$ with $\kappa \equiv \varphi_\pi \psi_4 + \varphi_y$; substituting this into the IS equation yields $y_t(1 - \psi_2 \kappa) = e_{1t} + \psi_2 \varphi_\pi e_{2t} + \psi_2 e_{3t}$. Back-substitution gives $Y_t = c_\theta + A_\theta e_t$ with

$$A_\theta = \frac{1}{D}\begin{pmatrix} 1 & \psi_2\varphi_\pi & \psi_2 \\ \psi_4 & 1 - \psi_2\varphi_y & \psi_2\psi_4 \\ \kappa & \varphi_\pi & 1 \end{pmatrix}, \qquad D = 1 - \psi_2\kappa = 1 - \psi_2(\varphi_\pi\psi_4 + \varphi_y). \tag{20}$$

Setting $\varphi_\pi = \varphi_y = 0$ recovers A_0 in (16). Since $\psi_2 < 0$ and $(\varphi_\pi, \varphi_y) \geq 0$, the determinacy denominator $D = 1 + |\psi_2|\kappa > 0$ for every rule on the grid: the static design has a unique stable solution everywhere, so no Taylor-principle restriction binds and the welfare-optimal rule may have $\varphi_\pi < 1$.

A.4 Variances and welfare

With $\Sigma_e = \mathrm{diag}(\sigma_1^2, \sigma_2^2, \sigma_3^2)$ the contemporaneous covariance is $\Sigma_Y = A_\theta \Sigma_e A_\theta'$, so

$$\mathrm{Var}(y) = \tfrac{1}{D^2}\left[\sigma_1^2 + \psi_2^2\varphi_\pi^2\sigma_2^2 + \psi_2^2\sigma_3^2\right], \tag{21}$$

$$\mathrm{Var}(\pi) = \tfrac{1}{D^2}\left[\psi_4^2\sigma_1^2 + (1-\psi_2\varphi_y)^2\sigma_2^2 + \psi_2^2\psi_4^2\sigma_3^2\right], \tag{22}$$

$$\mathrm{Var}(i) = \tfrac{1}{D^2}\left[\kappa^2\sigma_1^2 + \varphi_\pi^2\sigma_2^2 + \sigma_3^2\right], \tag{23}$$

and, the process being i.i.d., $\mathrm{Var}(\Delta i_t) = 2\ \mathrm{Var}(i_t)$. The welfare criterion (14) is evaluated from these closed forms.

B The medium-scale New Keynesian application

This appendix details the data, estimator, sampler, and policy experiment behind Section 6.2.

B.1 Data and VAR identification

We match the impulse responses of a reduced-form VAR on eight US quarterly series over 1965:Q3–1995:Q3: real GDP, real consumption, the GDP deflator, real investment, the real wage, labour productivity, the federal funds rate, and M2 growth.[2] The monetary policy shock is identified by a recursive (Cholesky) ordering with the funds rate seventh, so that it is orthogonal to the contemporaneous responses of the real block and the price level. M2 enters the VAR for identification only and is not a matching target; the matched responses are those of the first seven variables to the funds-rate shock.

B.2 Structural model and matching objective

The structural model is a medium-scale New Keynesian model following Christiano et al. (2005) in its main frictions: Calvo price and wage setting with indexation to lagged inflation, habit formation, investment adjustment costs, variable capital utilisation, and a working-capital cost channel, closed by an interest-rate rule. The sample-policy rule is part of the estimated environment, but the counterfactual policy experiment replaces its rule coefficients with the candidate action a. Eight parameters are estimated,

$$\theta = (\lambda_f,\ \xi_p,\ \xi_w,\ b,\ \kappa,\ \sigma_a,\ \sigma_m,\ \rho_r), \tag{24}$$

(price mark-up, Calvo price and wage stickiness, habit, investment adjustment cost, utilisation cost, monetary-shock standard deviation, and rule smoothing); the remaining pa-

[2]This drops real profits and the S&P 500 from the original ten-variable system in Christiano et al. (2005).

rameters are calibrated as in Christiano et al. (2005). Let $\hat{\gamma}_T$ stack the empirical responses of the seven matched variables to the monetary shock at horizons $j = 0, \dots, H-1$, and $\gamma(\theta)$ the model counterparts. The minimum-distance estimator $\hat{\theta}$ minimizes

$$Q_T(\theta) = \big[\hat{\gamma}_T - \gamma(\theta)\big]' \widehat{W} \big[\hat{\gamma}_T - \gamma(\theta)\big], \tag{25}$$

with a diagonal weight $\widehat{W} = \mathrm{diag}(1/\widehat{se}_j^2)$ from a residual bootstrap (Hall et al., 2012). Moments that the recursive ordering fixes at zero (the impact responses of the slow block) receive zero weight automatically.

B.3 Horizon selection

The matching horizon is chosen by the relative IRF-matching information criterion of Hall et al. (2012) (RIRSC), which selects $H = 7$ in the baseline empirical application.

B.4 Quasi-Bayes sampler and priors

Priors are uniform on $\lambda_f \in [1.001, 5]$, $\xi_p, \xi_w, b, \rho_r \in [0.01, 0.99]$, $\kappa \in [0.01, 20]$, $\sigma_a \in [0.001, 2]$, $\sigma_m \in [0.01, 5]$, so the quasi-posterior

$$\pi_{\mathrm{QB}}(\theta \mid X_T) \;\propto\; \exp\{-\tfrac{1}{2} Q_T(\theta)\}\, \pi(\theta) \tag{26}$$

is proportional to $\exp\{-\frac{1}{2}Q_T(\theta)\}$. We draw from it with a slice sampler (Neal, 2003), keeping 5,000 draws after 1,000 burn-in. For the policy step we thin to 500 draws. We first discard draws that fail the Blanchard–Kahn conditions at a small set of baseline Taylor-rule checks; in the welfare loop, unstable rule–draw combinations are recorded as missing. All 500 thinned draws pass the baseline filter in the reported run, and the baseline full-shock plug-in and quasi-Bayes rules are stable for all retained draws.

B.5 Policy experiment

The operational rule is the backward-looking Taylor rule

$$i_t = \rho_r\, i_{t-1} + (1-\rho_r)\big(\varphi_\pi \pi_{t-1} + \varphi_y y_{t-1}\big), \tag{27}$$

evaluated on a $16 \times 6 \times 7$ grid spanning $\varphi_\pi \in [0.5, 8]$, $\varphi_y \in [-0.5, 1.5]$, and $\rho_r \in [0, 0.99]$, for 672 rules. The grid is deliberately wider than the textbook range. For each rule and each draw, the unconditional variances entering $W(a, \psi) = -[\mathrm{Var}(\pi) + \mathrm{Var}(y) + 0.1\, \mathrm{Var}(\Delta i)]$ are computed from the model's state-space solution. The baseline welfare calculation uses the full model shock covariance matrix, including the calibrated auxiliary shocks needed to close the medium-scale model. Since the IRF-matching step identifies the model from the monetary-policy shock response, we also compute a monetary-shock-only welfare robustness check. Under this check the rule disagreement survives but becomes weak: $\widehat{a}_{\mathrm{QB}} = (3.00, 1.50, 0.00)$ and $\widehat{a}_{\mathrm{PI}} = (2.50, 1.50, 0.00)$, with a 0.21% posterior-mean welfare gain and 50.1% draw-level dominance. These rules are finite for 99.8% and 99.4% of retained draws, respectively, and the strict all-draw finite ranking gives a nearby quasi-Bayes rule. This is why the CEE exercise is interpreted qualitatively. The interest-rate-change term is the dynamic first difference variance, $\mathrm{Var}(\Delta i_t) = 2\,\mathrm{Var}(i_t) - 2\,\mathrm{Cov}(i_t, i_{t-1})$, not the static simplification used in Appendix A.4. The quasi-Bayes rule maximizes the quasi-posterior mean $\frac{1}{M}\sum_m W(a, \theta^{(m)})$ over draws; the plug-in rule maximizes $W(a, \hat{\theta})$.

References

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. and Inoue, A. (2003). The large sample behaviour of the generalized method of moments estimator in misspecified models. *Journal of Econometrics*, 114(2):361–394.

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.

Hansen, L. P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054.

Hansen, L. P., Heaton, J., and Yaron, A. (1996). Finite-sample properties of some alternative GMM estimators. *Journal of Business & Economic Statistics*, 14(3):262–280.

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

Neal, R. M. (2003). Slice sampling. *Annals of Statistics*, 31(3):705–767.

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.

Wald, A. (1950). *Statistical Decision Functions*. John Wiley & Sons, New York.

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton.

Woodford, M. (2010). Robustly optimal monetary policy with near-rational expectations. *American Economic Review*, 100(1):274–303.