

Description

This introduction estimates and interprets the parameters of a nonlinear New Keynesian model. We demonstrate the effect of a change in constraints on a model that is only partially identified.

Remarks and examples

Remarks are presented under the following headings:

The model

Parameter estimation

Policy and transition matrices

Impulse responses

A change in constraints

The model

Equations (1)–(5) specify a simplified nonlinear New Keynesian model of inflation Π_t , the output gap X_t , and the interest rate R_t . These three variables are driven by two state variables, Z_t and M_t . Recall that in this manual, capital letters denote the level of a variable, while lowercase letters denote deviations from steady state.

$$1 = \beta E_t \left(\frac{X_t}{X_{t+1}} \frac{1}{Z_t} \frac{R_t}{\Pi_{t+1}} \right) \quad (1)$$

$$(\theta - 1) + \phi(\pi_t - 1)\pi_t = \theta X_t + \phi \beta E_t \{(\pi_{t+1} - 1)\pi_{t+1}\} \quad (2)$$

$$\beta R_t = \Pi_t^\psi M_t \quad (3)$$

$$\ln(M_{t+1}) = \rho_m \ln(M_t) + u_{t+1} \quad (4)$$

$$\ln(Z_{t+1}) = \rho_z \ln(Z_t) + e_{t+1} \quad (5)$$

Equation (1) is the Euler equation linking the output gap in the current period to the expected future output gap and the real interest rate. In this equation, Z_t can be interpreted as a state variable influencing household decisions. In applications, Z_t stands in for the natural rate of interest, consumer spending shocks, or government expenditures. Equation (2) is the Phillips curve linking the current rate of inflation to expected future inflation and the output gap. Equation (3) is the interest-rate rule linking the interest rate to the inflation rate; in this equation, the state variable M_t captures all movements in the interest rate that are due to factors other than inflation. Equation (4) specifies the stochastic process for the state variable in the interest rate equation. Equation (5) specifies the stochastic process for the state variable in the Euler equation.

Parameter estimation

We estimate the model's parameters using data on the inflation rate and the nominal interest rate.

```
. use https://www.stata-press.com/data/r19/usmacro2
(Federal Reserve Economic Data - St. Louis Fed, 2017-01-15)
```

Not all model parameters are identified. We constrain β to be 0.96, a common value in the literature. The parameters θ and ϕ are collinear and cannot be estimated together. The parameter θ is set to 5, another common value in the literature. Later, we will explore consequences of the choice $\theta = 5$. The remaining six parameters are identified.

```
. constraint 1 _b[theta]=5
. constraint 2 _b[beta]=0.96
```

To specify this model to `dsgenl`, we will write each equation and then specify four options. The first three options tell `dsgenl` about the model's structure. The `observed()` option specifies observed control variables that are to be used in estimation. The `unobserved()` option specifies latent control variables. The `exostate()` option specifies the collection of state variables. The other option we use is the standard `constraints()` option that applies our constraints on the parameters β and θ .

```
. dsgenl (1 = {beta}*(x/F.x)*(1/z)*(r/F.p))
> ({theta}-1 + {phi}*(p -1)*p = {theta}*x + {phi}*{beta}*(F.p-1)*F.p)
> ({beta}*r = (p)~({psi=2})*m)
> (ln(F.m) = {rhom}*ln(m))
> (ln(F.z) = {rhoz}*ln(z)),
> exostate(z m) unobserved(x) observed(p r)
> constraint(1 2) nolog
Solving at initial parameter vector ...
Checking identification ...

First-order DSGE model
Sample: 1955q1 thru 2015q4                      Number of obs = 244
Log likelihood = -753.57131
( 1)  [/structural]theta = 5
( 2)  [/structural]beta = .96
```

	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
/structural						
beta	.96	(constrained)				
theta	5	(constrained)				
phi	47.07939	15.9416	2.95	0.003	15.83443	78.32435
psi	1.943008	.2957895	6.57	0.000	1.363271	2.522745
rhom	.7005489	.0452605	15.48	0.000	.61184	.7892577
rhoz	.9545255	.0186424	51.20	0.000	.9179871	.991064
sd(e.z)	.5689908	.0982979			.3763305	.761651
sd(e.m)	2.318208	.3047461			1.720916	2.915499

Because this is the nonlinear version of the model in [DSGE] Intro 3a, it is instructive to compare the two results. The two models deliver nearly identical estimates of most parameters; the small differences seen in the fourth and higher decimal places are due to differences in the numerical techniques used in each command. The only parameter that differs for the two models is the stickiness parameter ϕ , which is the counterpart of the Phillips curve slope κ from [DSGE] Intro 3a. The parameterization of the Phillips curve in (2) above differs from the one in [DSGE] Intro 3a. However, the two are linked by the relationship $\kappa = (\theta - 1)/\phi$, and we can use that relationship to show that the models are equivalent.

```
. nlcom (_b[theta] - 1)/_b[phi]
      _nl_1: (_b[theta] - 1)/_b[phi]
```

	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
_nl_1	.0849629	.0287694	2.95	0.003	.028576	.1413498

The value above is the same as the coefficient for kappa in [\[DSGE\] Intro 3a](#).

Policy and transition matrices

The policy matrix is part of the state-space representation of a DSGE model. It expresses the control variables of the model as functions of the state variables alone. Elements of the policy matrix represent the response of a control variable to a one-unit increase in a state variable. In a linear DSGE model, the policy matrix will also be linear. For a nonlinear DSGE model, the policy function will typically be nonlinear. The policy matrix reported after `dsge1` is a linear approximation of the true, nonlinear policy matrix evaluated at the estimated parameter vector. For details on the linearization of the policy function, see [DeJong and Dave \(2011, sec. 5.3\)](#).

```
. estat policy
```

Policy matrix

		Delta-method		z	P> z	[95% conf. interval]	
		Coefficient	std. err.				
x	z	.952921	.4813967	1.98	0.048	.0094009	1.896441
	m	-1.608216	.405057	-3.97	0.000	-2.402113	-.814319
p	z	.9678135	.277745	3.48	0.000	.4234433	1.512184
	m	-.4172515	.0393611	-10.60	0.000	-.494398	-.3401051
r	z	1.880469	.2615995	7.19	0.000	1.367744	2.393195
	m	.189277	.0591663	3.20	0.001	.0733133	.3052407

An increase in *m* decreases inflation, decreases the output gap, and increases interest rates. We may interpret it as an unexpected contractionary monetary policy. An increase in *z* increases inflation, increases the output gap, and increases interest rates.

Next, we can look at the state transition matrix. Because the states are uncorrelated with each other in this example, the elements of the state transition matrix are the persistence parameters in the model.

```
. estat transition
Transition matrix of state variables
```

		Delta-method		z	P> z	[95% conf. interval]	
		Coefficient	std. err.				
F.z	z	.9545255	.0186424	51.20	0.000	.9179871	.991064
	m	0	(omitted)				
F.m	z	0	(omitted)				
	m	.7005489	.0452605	15.48	0.000	.61184	.7892577

Note: Standard errors reported as missing for constrained transition matrix values.

These parameters repeat the results we saw in the estimation table. When the state variables are correlated, the transition matrix will reveal new information about the correlations among state variables.

Impulse responses

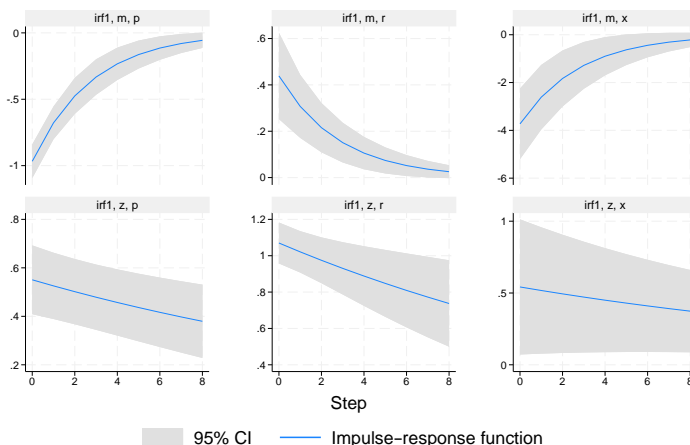
The impulse–response function traces out the effect of a shock on the model’s state and control variables. We use two commands to set up the impulse responses: an `irf set` command to set the active IRF file and an `irf create` command to create the impulse–response functions.

From there, we can use `irf graph irf` to graph the impulse–response functions.

```
. irf set nlexlirf, replace
(file nlexlirf.irf created)
(file nlexlirf.irf now active)

. irf create irf1, replace
(irfname irf1 not found in nlexlirf.irf)
(file nlexlirf.irf updated)

. irf graph irf, impulse(m z) response(x p r) byopts(yrescale)
```



Graphs by irfname, impulse variable, and response variable

Responses to a monetary shock are displayed in the first row. A shock to monetary policy leads inflation to fall, the interest rate to rise, and the output gap to fall. From the bottom row, we see that an increase in the natural rate of interest increases all three control variables.

A change in constraints

In the above example, the parameter θ was set to 5 to achieve identification of the other parameters. The following example shows the consequences of this decision by examining a case where θ was fixed to a different value. Suppose that instead of setting $\theta = 5$ above, we set $\theta = 2$. How would the change in constraints change our estimates? First, we store the current estimates with `estimates store`. We call the current estimates `theta5` because they represent the case where $\theta = 5$.

```
. dsge1
First-order DSGE model
Sample: 1955q1 thru 2015q4                      Number of obs = 244
Log likelihood = -753.57131
( 1)  [/structural]theta = 5
( 2)  [/structural]beta = .96
```

	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
/structural						
beta	.96	(constrained)				
theta	5	(constrained)				
phi	47.07939	15.9416	2.95	0.003	15.83443	78.32435
psi	1.943008	.2957895	6.57	0.000	1.363271	2.522745
rhom	.7005489	.0452605	15.48	0.000	.61184	.7892577
rhoz	.9545255	.0186424	51.20	0.000	.9179871	.991064
sd(e.z)	.5689908	.0982979			.3763305	.761651
sd(e.m)	2.318208	.3047461			1.720916	2.915499

```
. estimates store theta5
```

We can now recall these estimates later as needed. We set up the new constraint with

```
. constraint 3 _b[theta]=2
```

and reestimate parameters. The `nolog` option is used to suppress the iteration log.

```
. dsge1 (1 = {beta}*(x/F.x)*(1/z)*(r/F.p))
> ({theta}-1 + {phi}*(p -1)*p = {theta}*x + {phi}*{beta}*(F.p-1)*F.p)
> (({beta})*r = (p)~({psi=2})*m)
> (ln(F.m) = {rhom}*ln(m))
> (ln(F.z) = {rhoz}*ln(z)),
> exostate(z m) unobserved(x) observed(p r)
> constraint(2 3) nolog
Solving at initial parameter vector ...
Checking identification ...

First-order DSGE model
Sample: 1955q1 thru 2015q4                      Number of obs = 244
Log likelihood = -753.57131
( 1) [/structural]beta = .96
( 2) [/structural]theta = 2
```

	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
/structural						
beta	.96	(constrained)				
theta	2	(constrained)				
phi	11.76979	3.985337	2.95	0.003	3.958675	19.58091
psi	1.943004	.2957868	6.57	0.000	1.363273	2.522736
rhom	.7005482	.04526	15.48	0.000	.6118401	.7892562
rhoz	.9545256	.0186424	51.20	0.000	.9179872	.991064
sd(e.z)	.5689892	.0982974			.37633	.7616485
sd(e.m)	2.318204	.3047434			1.720918	2.91549

```
. estimates store theta2
```

We saved these estimates in `theta2` because they represent the case where $\theta = 2$. `estimates table` displays the two results. The `stats(11)` option displays the estimated log likelihood with the estimated structural parameters.

```
. estimates table theta5 theta2, b(%9.4f) stats(11)
```

Variable	theta5	theta2
/structural		
beta	0.9600	0.9600
theta	5.0000	2.0000
phi	47.0794	11.7698
psi	1.9430	1.9430
rhom	0.7005	0.7005
rhoz	0.9545	0.9545
sd(e.z)	0.5690	0.5690
sd(e.m)	2.3182	2.3182
Statistics		
11	-753.5713	-753.5713

The reported log-likelihood values are identical across the two parameterizations. The change in constraints slides us along the top of a ridge in the likelihood function. The estimated value of `phi` differs when the constraint is changed, but all other parameter estimates are identical. Even though the point estimate of `phi` has changed, the z statistic and p -value are unchanged.

As discussed above, the underlying parameter that is being estimated here is $\kappa = (\theta - 1)/\phi$. The restrictions on the parameter θ lead to a change in the estimated value of ϕ but do not change the estimated value of κ . We can use `nlcom` to recover κ first from the model with $\theta = 5$ and then from the model with $\theta = 2$.

```
. estimates restore theta5
(results theta5 are active now)
. nlcom (_b[theta]-1) / _b[phi]
      _nl_1: (_b[theta]-1) / _b[phi]
```

	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
_nl_1	.0849629	.0287694	2.95	0.003	.028576	.1413498

```
. estimates restore theta2
(results theta2 are active now)
. nlcom (_b[theta]-1) / _b[phi]
      _nl_1: (_b[theta]-1) / _b[phi]
```

	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
_nl_1	.0849633	.0287692	2.95	0.003	.0285767	.1413498

The equivalence of the two parameterizations can also be seen through the policy matrix. The estimated policy matrix when $\theta = 2$ is

```
. estat policy
Policy matrix
```

		Delta-method					
		Coefficient	std. err.	z	P> z	[95% conf. interval]	
x	z	.9529193	.481392	1.98	0.048	.0094082	1.89643
	m	-1.608214	.4050521	-3.97	0.000	-2.402102	-.8143265
p	z	.9678173	.2777448	3.48	0.000	.4234476	1.512187
	m	-.4172521	.0393609	-10.60	0.000	-.4943981	-.3401062
r	z	1.880473	.2615995	7.19	0.000	1.367748	2.393199
	m	.1892774	.059166	3.20	0.001	.0733142	.3052406

This policy matrix is identical to the one shown above, in which the constraint $\theta = 5$ was applied. The two parameter vectors imply the same reduced-form matrices.

Reference

DeJong, D. N., and C. Dave. 2011. *Structural Macroeconometrics*. 2nd ed. Princeton, NJ: Princeton University Press.
<https://doi.org/10.2307/j.ctt7srm7>.

Also see

[DSGE] [Intro 1](#) — Introduction to DSGEs

[DSGE] [Intro 3a](#) — New Keynesian model

[DSGE] [Intro 3e](#) — Nonlinear New Classical model

[DSGE] [Intro 3f](#) — Stochastic growth model

[DSGE] [dsgenl](#) — Nonlinear dynamic stochastic general equilibrium models

[DSGE] [dsgenl postestimation](#) — Postestimation tools for dsgenl

Stata, Stata Press, and Mata are registered trademarks of StataCorp LLC. Stata and Stata Press are registered trademarks with the World Intellectual Property Organization of the United Nations. StataNow and NetCourseNow are trademarks of StataCorp LLC. Other brand and product names are registered trademarks or trademarks of their respective companies. Copyright © 1985–2025 StataCorp LLC, College Station, TX, USA. All rights reserved.

For suggested citations, see the FAQ on [citing Stata documentation](#).

