The Half-Credible Economy

What forward-looking expectations change in the engineered monetary framework, why partial credibility is more dangerous than none, and the correction it forces on the rule

GFE-PUB-EMF-20261031-A · 31 October 2026
Dr. Glen Brown · Global Internal Governance Chamber · Global Financial Engineering, Inc.
Public edition.

Abstract

The two earlier papers of this series, The Better Ways and Room Above the Floor, built and calibrated an engineered monetary framework inside a model in which people form their expectations of inflation and output from the past. Such a model cannot see the most important thing a central bank does: shape what people expect of the future. I said so in both papers, and promised the test. This paper is that test. I give households and firms a weight on the model-consistent future, from none to most, solve the economy under rational expectations with the interest-rate floor enforced as an anticipated constraint, re-tune every rule at every weight, and run the five tests I wrote down before any result was computed.

The framework survives, with one correction. Once that correction is made, at every forward weight, at every neutral nominal rate from 3 to 6 per cent, and under both fixed and variable mortgages, the engineered rule beats the best conventional rule tuned on the same data, by 3 to 17 per cent at a neutral nominal rate of four. Its two contested design choices, that the rate should respond to supply-driven inflation rather than look through it, and that the supply and credit instruments should stand down near the floor, hold as well. But the test found two things the backward model had hidden. The first is that credibility is not a straight line. When expectations are largely forward-looking, the floor almost ceases to matter. When they are only partly forward-looking, it is more dangerous than in the backward model: at a neutral nominal rate of three, the share of paths that fall into a deflation trap rises from 2 to 9 per cent. The rule from Room Above the Floor that four points of neutral nominal rate is an upper estimate of the room an economy needs is therefore wrong where credibility is partial, and I correct it here. The second is a defect in the rule itself: its easing for mortgage stress, unbounded, can pull the policy rate's response to inflation below one for one, which a backward economy tolerates and a forward-looking one does not. Bounding it at one point removes the failure. The bounded rule is EMF-4.

I. The test the earlier papers owed

Every framework tested in a backward-looking model is open to one objection, and it is the oldest in modern macroeconomics. Lucas (1976) argued that the relationships such models estimate are not fixed features of the economy but responses to the policy in force; change the policy, and people who understand it change their behaviour. The objection bites hardest on exactly the questions this series has asked. How much room a central bank needs above the floor depends on whether people expect it to keep rates low once the floor is behind it. Whether a rate rule should look through a supply shock depends on whether people believe the shock will pass. Whether easing for mortgage stress is safe depends on whether people read it as a lapse in the commitment to price stability.

The standard answer is the hybrid New Keynesian economy, in which inflation and output depend partly on their own past and partly on what people expect of the next quarter, and in which those expectations are consistent with the model and the policy rule. Galí and Gertler (1999) estimated that forward-looking behaviour dominates in American inflation, though later work finds a wide range depending on the data and the method. Rather than choose one estimate, I test three weights: none, which reproduces the earlier papers; 0.3, an economy in which people look ahead somewhat but are still steered mainly by recent experience; and 0.6, one in which the future dominates. I call the weight credibility in what follows, because in this model it measures how far the public's expectations are formed from the central bank's expected conduct rather than from the past.

II. The forward-looking economy

The economy is that of Room Above the Floor, with one change: in the Phillips curve, the spending equation and the real interest rate that enters both, a share λ of the backward-looking terms is replaced by the expectation of next quarter's value. With λ equal to zero the economy is exactly the earlier one; I confirmed that the new code reproduces the earlier results to the fourteenth decimal place, and the re-run of the earlier experiments below matches the published figures. The linear part of the economy is solved under rational expectations. The policy-rate floor of −1 per cent is enforced as an anticipated constraint, using the method of Holden and Paetz (2012): when the rate would breach the floor, the solution adds the smallest set of anticipated future rate adjustments that keeps the whole expected path at or above it, so that people know the floor will bind and for how long. The calculation is repeated every quarter. The nonlinear parts of the engineered rule (the easing for mortgage stress, the credit instrument, supply relief and the floor gate) are computed each quarter and enter as policy actions that people observe but do not anticipate beyond the quarter.

No rule is judged with coefficients tuned for another economy. At each forward weight I re-tuned the best conventional rule, the engineered rule (EMF-3 of Room Above the Floor) and a version of it that looks through supply inflation, on twelve training economies and training shocks, and then evaluated every rule on the fifty-four test economies and unseen shocks of the earlier papers, at neutral nominal rates from 1.5 to 6 per cent and under long fixed and variable mortgages. The loss is the household-weighted loss of the earlier papers.

III. Five tests, and two departures from the plan

Before computing any forward-looking result I wrote down five tests and the condition under which each would count as passed or reversed. Table 1 gives the verdicts. Two departures from the plan were forced by what the run found, and I report both before the results so that the reader can weigh them.

The first concerns the loss. At a forward weight of 0.3, some paths fall into a deflation trap: the rate sits at the floor, expected deflation raises the real interest rate, and inflation falls without limit. Such a path has an unbounded loss, and a handful of them dominated the mean so completely that the first round of tuning produced degenerate rules. The trap is genuine model behaviour: doubling the horizon of the floor calculation from 24 to 48 quarters left every trapped path unchanged. I therefore capped the loss of each path at 50, about ten times that of a typical path, so that a trap counts as a heavy but finite penalty, re-tuned the rules at that weight only, and report the share of trapped paths as an outcome in its own right. At the

other two weights no path reaches the cap, so their results are unaffected. Because the cap applies everywhere, the losses at low room in this paper are smaller than the uncapped means reported in Room Above the Floor, though the ordering of rules is the same.

The second concerns the rule. At a forward weight of 0.3 and a neutral nominal rate of six, the engineered rule lost to the conventional one, and some of its paths exploded upward. Section VI traces the cause to the rule's easing for mortgage stress, and tests a remedy fixed once, at a value chosen before the test and not tuned: a bound of one point on how far that easing may lower the policy rate.

Table 1. Tests fixed in advance and their verdicts

Test	Criterion written before the run	Verdict
R1 Framework	EMF-3 beats the best conventional rule at a neutral nominal rate of 4, at forward weights 0.3 and 0.6, with fixed and with variable mortgages	Holds in all four cases: by 5% and 15% at 0.3, by 3% and 4% at 0.6
R2 Room	The rule of four is confirmed as an upper estimate if the best neutral nominal rate at weight 0.6 is 4 or below	Holds at 0.6, where outcomes are flat from 1.5 to 5. But at 0.3 the best rate rises to 5 with fixed mortgages: four is not an upper estimate there (Section IV)
R3 Look-through	The earlier finding that the rate should respond to supply inflation reverses if a rule that looks through it does better	Does not reverse: looking through is worse by 16–18% at 0.3 and by 9% at 0.6
R4 Gate	EMF-3 is no worse than the ungated EMF-2 at a neutral nominal rate of 3	Holds: better by 5% at 0.3; identical at 0.6, where the floor is not reached
R5 Floor	Time at the floor and loss at a neutral nominal rate of 3 fall as the forward weight rises	Fails: both rise from weight 0 to 0.3 before falling to almost nothing at 0.6

IV. Credibility is not a straight line

The failure of the fifth test is the central result of this paper, and it was not what I expected. The conventional view, which I repeated in Room Above the Floor, is that forward-looking expectations make the floor less costly: if people know the central bank will keep rates low after the floor stops binding, expected inflation rises today, the real rate falls, and part of the work the rate cannot do is done by the promise (Eggertsson and Woodford, 2003; Reifschneider and Williams, 2000). At a forward weight of 0.6 the model agrees, and emphatically. Shocks fade so quickly that the policy rate rarely needs to move far; the floor binds in at most 1.5 per cent of quarters even when the neutral nominal rate is only 1.5; no path is ever trapped; and inflation error is half what it is in the backward economy.

At a forward weight of 0.3 the opposite happens. Table 2 and Figure 1 show that at every level of room the loss is higher than in the backward economy; up to a neutral nominal rate of four the floor binds more often; and wherever traps occur, the share of paths that fall into one is between two and eleven times larger. At a neutral nominal rate of three, 9 per cent of paths are trapped against 2 per cent in the backward economy; at 2.5, 17 per cent against 4. [Inference] The mechanism is the one Benhabib, Schmitt-Grohé and Uribe (2001) identified in a different setting. In a partly forward-looking economy, expected deflation already raises the real interest rate at the floor, which is the channel through which expectations hurt; but expectations are not yet forward-looking enough for the anticipated return of policy to pull inflation back, which is the channel through which they help. The economy has acquired the vulnerability of a credible regime without its protection.

Table 2. Household-weighted loss and deflation traps by neutral nominal rate and forward weight, long fixed mortgages, 54 economies

Neutral nominal rate (r* + target)	1.5	2	2.5	3	3.5	4	5	6
Weight 0: EMF-3	5.96	5.02	4.45	4.06	3.91	3.85	3.88	4.56
Weight 0: best conventional	5.98	5.10	4.58	4.25	4.14	4.10	4.28	5.26
Weight 0: trapped paths (%)	17.2	9.1	4.1	1.9	0.7	0.2	0.0	0.0
Weight 0.3: EMF-3	**18.28**	**13.38**	**9.97**	**7.45**	**5.90**	**5.32**	**4.70**	**7.39**
Weight 0.3: best conventional	18.34	13.27	10.15	7.65	6.17	5.59	5.26	5.99
Weight 0.3: trapped paths (%)	**45.5**	**29.4**	**17.4**	**9.3**	**4.5**	**2.2**	**0.3**	**0.0**
Weight 0.6: EMF-3	1.68	1.68	1.68	1.68	1.68	1.68	1.68	2.00
Weight 0.6: best conventional	1.74	1.74	1.74	1.74	1.74	1.74	1.74	2.34
Weight 0.6: trapped paths (%)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

[Derived] Each path's loss is capped at 50, so trapped paths count as a heavy finite penalty. A path is trapped if four-quarter inflation ever falls below −4 per cent. The rise of EMF-3 at a neutral nominal rate of 6 under weight 0.3 is the defect examined in Section VI.

Figure 1. Loss and deflation traps by neutral nominal rate at three forward weights

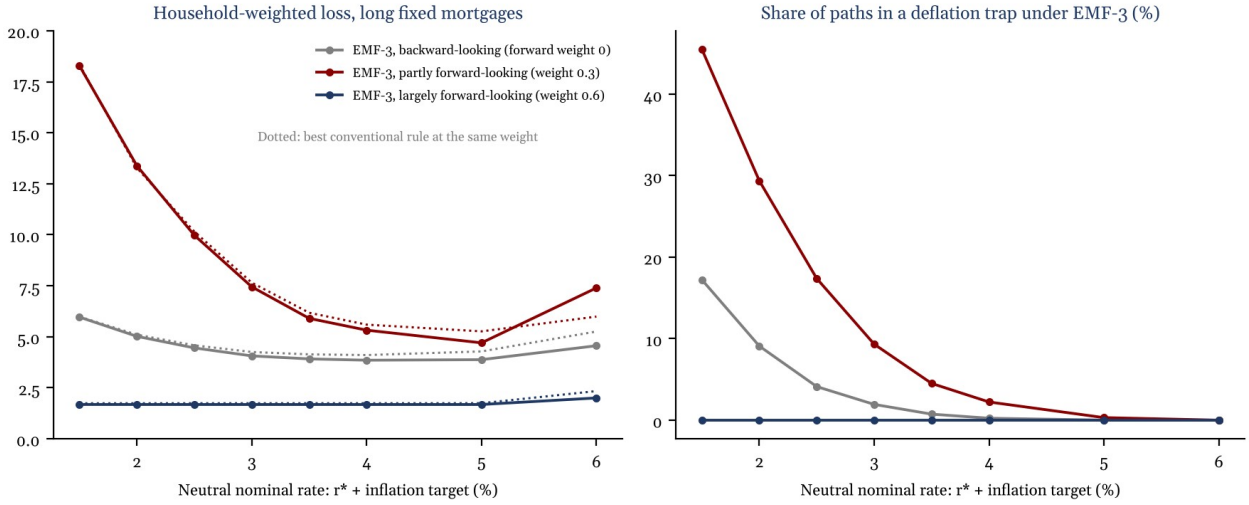

[Derived] Largely forward-looking expectations almost abolish the floor problem. Partly forward-looking expectations make it worse than no forward-looking expectations at all.

What this does to the rule of four

Room Above the Floor concluded that outcomes are best when the neutral nominal rate, the neutral real rate plus the inflation target, is about four, and described that figure as an upper estimate, on the ground that forward-looking expectations would reduce the room needed. That reasoning holds only where credibility is high. At a forward weight of 0.6, every neutral nominal rate from 1.5 to 5 gives practically the same result, and four is indeed generous. At a forward weight of 0.3, the best result with long fixed mortgages comes at five, with a loss about 12 per cent below that at four; with variable mortgages, under the bounded rule of Section VI, four and five tie, because the higher nominal rate raises debt service as much as the extra room saves. Since no central bank knows with confidence where its public's expectations lie between these cases, the robust choice is the one that is safe in the worst of them.

I therefore revise the rule. The neutral nominal rate should be four to five, not about four: five where mortgages are long and fixed and credibility is not yet secure, four where mortgages float. The inflation target follows as four to five minus the best estimate of the neutral real rate. Where that rate is 2 per cent, the target remains 2 to 3; where it is 1 per cent, 3 to 4. [Inference] A central bank that has earned high credibility can sit at the lower end, but it should not assume that it has, because the cost of being wrong falls entirely on the side of too little room.

V. What survives

Table 3 sets out the comparisons on which the first, third and fourth tests rest. The engineered rule beats the best conventional rule at every forward weight. Its advantage is largest where mortgages float and expectations look backward, 21 per cent, because that is where its easing for mortgage stress does most; it narrows to 3 to 4 per cent at a forward weight of 0.6, because there the policy rate moves so little that debt service rarely approaches the stress threshold. The conventional benchmark matters too: the textbook Taylor rule, not tuned at all, does twice as badly as either tuned rule at a forward weight of 0.3, because it is too slow to keep the economy away from the floor.

Table 3. Household-weighted loss at a neutral nominal rate of 4, and the gate at 3

Rule	Weight 0 fixed	Weight 0 variable	Weight 0.3 fixed	Weight 0.3 variable	Weight 0.6 fixed	Weight 0.6 variable
EMF-3	**3.85**	**4.19**	**5.32**	**5.57**	**1.68**	**1.68**
Best conventional rule	4.10	5.29	5.59	6.56	1.74	1.75
EMF-3 looking through supply inflation	4.00	4.44	6.19	6.56	1.84	1.84
Textbook Taylor rule	5.42	6.40	12.84	13.97	1.83	1.84
At a neutral nominal rate of 3: EMF-3 (gated)	4.06	4.17	7.45	7.32	1.68	1.68
At a neutral nominal rate of 3: EMF-2 (no gate)	4.17	4.27	7.85	7.68	1.68	1.68

[Derived] Every rule except the textbook Taylor rule is re-tuned at each forward weight. The look-through rule has a zero response to supply inflation and is tuned on the same data.

The look-through result deserves a sentence, because forward-looking expectations were the strongest case against it. The usual argument for looking through a supply shock is that it will pass of its own accord, and that a credible central bank can wait for it to do so without unsettling expectations. In this model the argument fails at every weight. Looking through is worse by 16 to 18 per cent at a weight of 0.3 and by 9 per cent at 0.6. [Inference] The reason is that supply shocks here are persistent and partly misread, as they are in practice, and a central bank that responds to none of them invites the public to discount its response to all of them. The gate, which switches the supply and credit instruments off as the rate nears the floor, is likewise vindicated where the floor matters and costs nothing where it does not.

VI. The flaw the backward model could not see

The engineered rule lowers the policy rate when households' debt service exceeds 10 per cent of income, in proportion to the excess. In The Better Ways this easing was one of the rule's main sources of advantage for households, and it remains so. But it carries a hidden cost. Once mortgage rates have caught up with the

policy rate, every point by which the rate rises adds a point to debt service, and the easing takes part of that point back. If the rule moves the rate by 1 + a for each point of inflation and eases by b for each point of excess debt service, then while debt service is above the threshold the rate's long-run response to inflation is (1 + a) / (1 + b). When that ratio is below one, the rule raises the nominal rate by less than inflation, so the real rate falls as inflation rises.

In a backward-looking economy whose expectations are anchored to the target, that is harmless: the anchor, not the rate, brings inflation home, and the rule tuned in The Better Ways had a ratio of 0.53 without ill effect. In a forward-looking economy it is not harmless, because there the requirement that the real rate rise with inflation, known as the Taylor principle, is what holds expectations in place. Re-tuned at a forward weight of 0.3, the rule's coefficient on mortgage stress rose to 3.1, and the ratio fell to 0.40. In the training economies, with a neutral nominal rate of four, baseline debt service sits below the threshold, the easing is rarely active, and the weakness never showed. At a neutral nominal rate of six, baseline debt service sits above the threshold, the easing is permanently engaged, and some paths spiralled upward. [Derived] The arithmetic is exact; [Inference] that it is the cause is shown by the remedy.

The remedy is a bound. However large the excess of debt service, the easing may lower the policy rate by no more than one point. Beyond that, each further point of inflation moves the rate by the full 1 + a, and the Taylor principle is restored where it matters, in large deviations. I set the bound at one point before testing it and did not tune it, and I did not re-tune the rule's coefficients around it. Table 4 and Figure 2 show the result. At a forward weight of 0.3 and a neutral nominal rate of six, the bound cuts the loss from 12.8 to 6.8 with variable mortgages and from 7.4 to 5.0 with fixed ones, so that the rule again beats the conventional benchmark. It has a price: in the backward economy with variable mortgages, where unbounded easing was safe, it raises the loss by 2 to 8 per cent, and the rule's advantage at a neutral nominal rate of four narrows from 21 to 17 per cent. With the bound, the engineered rule beats the best conventional rule in all twenty-four combinations of forward weight, neutral rate from three to six, and mortgage type. That is the trade I recommend, because the bound costs a little where the old rule was safe and saves a great deal where it was not.

Table 4. The bound on mortgage easing: household-weighted loss by neutral nominal rate

Economy and rule	3	4	5	6
Weight 0, variable: EMF-3 (unbounded)	4.17	4.19	4.74	6.21
Weight 0, variable: EMF-4 (bounded)	**4.26**	**4.40**	**5.12**	**6.65**
Weight 0, variable: best conventional	4.80	5.29	6.50	8.63
Weight 0.3, variable: EMF-3 (unbounded)	7.32	5.57	7.14	12.81
Weight 0.3, variable: EMF-4 (bounded)	**7.35**	**5.51**	**5.50**	**6.79**
Weight 0.3, variable: best conventional	8.03	6.56	7.28	9.24
Weight 0.3, fixed: EMF-3 (unbounded)	7.45	5.32	4.70	7.39
Weight 0.3, fixed: EMF-4 (bounded)	**7.45**	**5.32**	**4.65**	**5.04**
Weight 0.3, fixed: best conventional	7.65	5.59	5.26	5.99
Weight 0.6, variable: EMF-4 (bounded)	1.68	1.68	1.74	2.28
Weight 0.6, variable: best conventional	1.74	1.75	1.90	2.82

[Derived] EMF-4 is EMF-3 with the easing for mortgage stress bounded at one point of the policy rate, coefficients unchanged. At weight 0.6 the bound changes little because debt service rarely reaches the threshold.

Figure 2. The bound on mortgage easing, and the framework's advantage at every forward weight

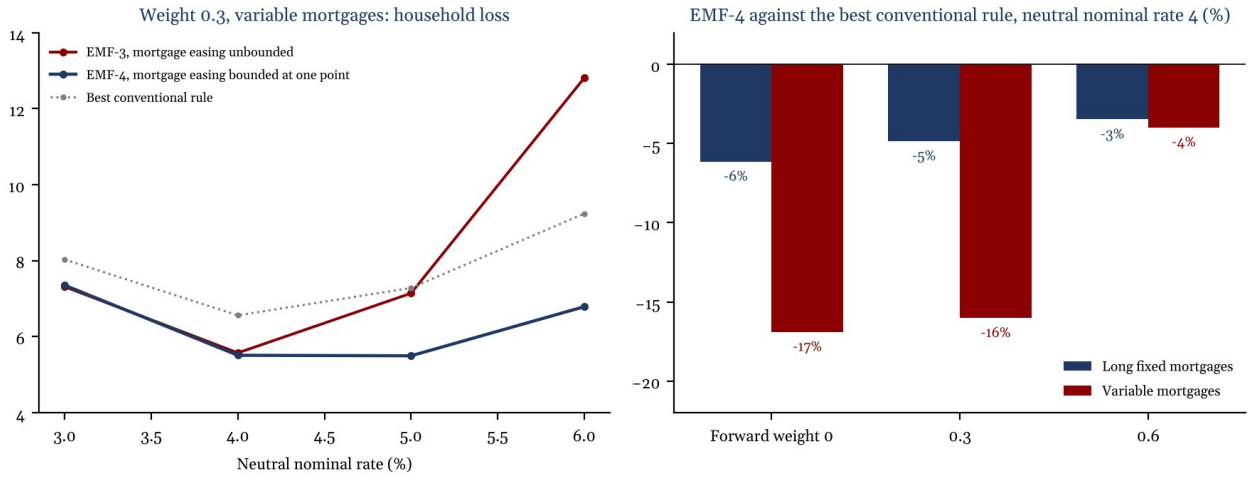

[Derived] Left: unbounded easing fails where baseline debt service exceeds the threshold; the bound removes the failure. Right: with the bound, the framework beats the best conventional rule at every forward weight.

VII. Corrections to the earlier papers

A series that tests its own claims must say plainly where they did not survive. Three corrections follow from this paper, and the web editions of the earlier papers now carry a note pointing here.

Earlier paper and claim	Correction
Room Above the Floor: the rule of four is an upper estimate of the room an economy needs, because forward-looking expectations reduce the cost of the floor.	True only where credibility is high. Where it is partial, the floor is more dangerous than in the backward model and the best neutral nominal rate rises to about five with fixed mortgages. The rule becomes four to five (Section IV).
The Better Ways: the rate rule eases in proportion to excess mortgage debt service, without limit.	The easing must be bounded, at one point, so that the rate's response to inflation is never pulled below one for one. The bounded rule is EMF-4 (Section VI).
Room Above the Floor: losses at very low room are reported as means across paths.	Those means were dominated by trapped paths. With each path's loss capped, the ordering of rules is unchanged but the magnitudes are much smaller; the share of trapped paths is the better measure (Section III).

Everything else in the two earlier papers stands as written: the rule's credit instrument and supply relief, the floor gate, the response to supply inflation, the preference for a point target over a band, the mortgage tilt and its removal by indexation, and the design of the reset cap. The cap was not re-tested under forward-looking expectations, and I make no new claim about it.

VIII. Guidance

The framework that emerges from the three papers, revised by this one, can be stated as eight rules.

Rule	Basis
1. Set the inflation target from the room it buys: aim for a neutral nominal rate of four to five, so the target is four to five minus the best estimate of the neutral real rate. Use five where mortgages are long and fixed and credibility is not secure; four where mortgages float.	Table 2; Room Above the Floor
2. Restore room, but do not go beyond it. Moving to the range is worth a great deal; moving past it is worth almost nothing.	Room Above the Floor
3. Treat credibility as an instrument and build it deliberately. Partial credibility is the most dangerous state: it brings the deflationary pull of expectations without the anchoring that full credibility provides. Announce the rule, explain it, and keep to it.	Section IV; inference
4. Run the policy rate by the engineered rule, responding to supply-driven inflation rather than looking through it.	Table 3; The Better Ways
5. Ease for mortgage stress, but by no more than one point of the policy rate, so that the rate always rises by more than inflation.	Table 4
6. Use the credit instrument and fast supply relief, and stand both down as the rate approaches the floor.	Table 3; Room Above the Floor
7. Prefer a point target, reviewed on a fixed schedule against new estimates of the neutral real rate.	Room Above the Floor
8. Cap upward mortgage resets by deferral repaid on a schedule, or by insurance funded outside the banks; never let the cost fall on bank capital.	Room Above the Floor

The complaint that began this series was that monetary policy is run to hit a number at the expense of the people who live with the consequences. The forward-looking test sharpens the answer rather than overturning it. The engineered framework serves households better than the best conventional rule whether the public looks back or ahead. What the test adds is that the public's view of the future is itself part of the machinery, and the most fragile setting of that machinery is the middle one. [Inference] A central bank that is believed only partly is exposed in a way that neither a fully trusted nor a wholly ignored one is, and the remedy is not a cleverer rule but a clearer one, kept long enough to be believed.

IX. Limits

The forward weight is a single number applied to every household and firm, and it is chosen rather than estimated; real economies mix agents who look ahead with agents who do not, and the mix may change with the regime. The forward-looking expectations are fully model-consistent, which overstates what anyone actually knows. The engineered rule's nonlinear instruments are treated as surprises beyond the current quarter, so the model does not credit the rule with the extra stabilisation that anticipating them might bring. The shocks have the same size at every weight; a forward-looking economy might in practice face larger or different shocks, and the strikingly calm results at a weight of 0.6 should be read with that in mind. The model contains no fiscal policy, which in real episodes at the floor has done much of the work, so the trap rates at a weight of 0.3 describe an economy left to its central bank alone. Two choices were made after seeing results and are reported as such: the cap on each path's loss, and the bound on mortgage easing, which was set once at one point and not tuned. Every number here is model-dependent; the orderings are more robust than the magnitudes.

Table 5. The seam: principal claims by evidential status

Claim	Status
The engineered rule beats the best conventional rule at every forward weight tested	Derived
Partly forward-looking expectations make the floor more dangerous than backward-looking ones	Derived
The mechanism is expected deflation raising the real rate without the pull of anticipated policy	Inference
The robust neutral nominal rate is four to five; the target is four to five minus the neutral real rate	Derived; inference as policy
Unbounded easing for mortgage stress can break the Taylor principle; a one-point bound restores it	Derived
Responding to supply inflation beats looking through it under forward-looking expectations	Derived
Forward-looking behaviour is substantial in measured inflation dynamics	Witnessed in the literature

Method note

The model replaces a share λ of the backward-looking terms of the earlier model with model-consistent expectations of next quarter, and reduces exactly to the earlier model at λ = 0. The linear part is solved by iterating on the rational-expectations solution; the floor is imposed by anticipated shadow-rate adjustments found each quarter by solving a linear complementarity problem over a 24-quarter horizon, and the results do not change at 48 quarters. Rules are re-tuned at each weight on twelve training economies with eighty training paths, at a neutral nominal rate of four with long fixed mortgages, and evaluated on fifty-four economies with two hundred unseen paths of 160 quarters after a burn-in of forty. The five tests were written down before any forward-looking result was computed. The cap on each path's loss and the bound on mortgage easing were introduced after the first results and are reported as departures from the plan.

References: Benhabib, J., S. Schmitt-Grohé and M. Uribe (2001), "The Perils of Taylor Rules", Journal of Economic Theory 96. Eggertsson, G. and M. Woodford (2003), "The Zero Bound on Interest Rates and Optimal Monetary Policy", Brookings Papers on Economic Activity 2003(1). Galí, J. and M. Gertler (1999), "Inflation Dynamics: A Structural Econometric Analysis", Journal of Monetary Economics 44. Holden, T. and M. Paetz (2012), "Efficient Simulation of DSGE Models with Inequality Constraints", University of Surrey Discussion Paper 16/12. Lucas, R. (1976), "Econometric Policy Evaluation: A Critique", Carnegie-Rochester Conference Series on Public Policy 1. Reifschneider, D. and J. Williams (2000), "Three Lessons for Monetary Policy in a Low-Inflation Era", Journal of Money, Credit and Banking 32. Rudebusch, G. and L. Svensson (1999), "Policy Rules for Inflation Targeting", in J. Taylor (ed.), Monetary Policy Rules, University of Chicago Press.

Paper of the Public Doctrine Series, Global Financial Engineering, Inc. It is analysis, not advice, and it takes no position on any particular decision of any central bank or government.