PITFALLS OF LINEAR APPROXIMATION OF STOCHASTIC MODELS

1. A LIST OF PITFALLS

- Linearized models are of course valid only locally. In stochastic economic models, this usually means valid for small stochastic disturbances that don't push us far from the non-stochastic steady state. But the restrictions this places on the validity of linearization are sometimes overdrawn. For example, linearization can be justified, for certain purposes, in non-stationary models.
- There is one apparently natural interpretation of "replacing the model with a linear-quadratic approximation" that does not lead to valid first-order expansions.
- Simulating the first-order approximation to a model and computing average welfare from the simulations does not generally lead to correct rankings of welfare across alternative policies. Higher than first order approximation is generally needed.

2. LINEARIZATION WITH NONSTATIONARITY AND UNBOUNDED DISTURBANCES

Suppose we have a dynamic system

$$x_{t+1} = f(x_t, \varepsilon_{t+1}) \,. \tag{1}$$

We can apply it recursively to obtain

$$x_{t+1} = f(f(x_{t-1}, \varepsilon_t), \varepsilon_{t+1}) \,. \tag{2}$$

Assuming the usual conditions are met, we can take a Taylor expansion of it in $\varepsilon_t, \varepsilon_{t+1}$ in the neighborhood of $\varepsilon_t = \varepsilon_{t+1} = 0$ and $x_{t-1} = x^*$ by applying the chain rule for differentiating functions of functions:

$$\begin{aligned} x_{t+1} - f(f(x^*,0),0) \doteq & \\ & D_1 f(f(x^*,0),0) D_2 f(x^*,0)\varepsilon_t + D_2 f(f(x^*),0)\varepsilon_{t+1} \\ & \qquad + D_1 f(f(x^*,0),0) D_1 f(x^*,0)(x_{t-1} - x^*) \,. \end{aligned} \tag{3}$$

This expansion will be, like any Taylor expansion, accurate for small enough ε_t and ε_{t+1}. The error of approximation will grow small relative to the size of the ε's as the ε's get small.

If $x^* = \bar{x}$, where $\bar{x} = f(\bar{x}, 0)$ (i.e. $\bar{x}$ is a deterministic steady state value for x_t), We can apply these ideas recursively to get the s-step Taylor expansion

$$x_t - \bar{x} = \sum_{s=0}^{t-1} D_1 f(\bar{x}, 0)^s D_2 f(\bar{x}, 0) \varepsilon_{t-s} + D_1 f(\bar{x}, 0)^t (x_0 - \bar{x}) \,. \tag{4}$$

3

- This formula does require $x^* = \bar{x}$. Otherwise, we would obtain a much more complicated expression involving derivatives at many different values of x. Such a "linearization about a deterministic path" can sometimes be useful, but it is not common in macroeconomics.
- The validity of the Taylor expansion for small ε and small deviations of x_0 from $\bar{x}$ does not depend on any condition on the eigenvalues of $D_1 f(\bar{x}, 0)$.
- If we want to take a limit as $t \to \infty$ and still get a valid approximation, we do require that all the eigenvalues of $D_1 f(\bar{x}, 0)$ be less than one in absolute value, as otherwise the limiting expression would not be defined. If this condition is met, and if we can uniformly bound all the ε_t's, we can make Taylor-series-approximation accuracy claims as the bound on the size of the ε's shrinks in this $t \to \infty$ case.

4

- If the ε's are random and have unbounded support, e.g. i.i.d. $N(0, \sigma^2)$, we will not be able to make non-probabilistic accuracy claims. There will always be some probability that ε's are large enough to invalidate the Taylor approximation, no matter how small we make σ.
- In this case we also cannot let $t \to \infty$ and make accuracy claims. The probability that *eventually* some ε will be large enough to invalidate the approximation is one, no matter how small is λ. Once we are out of the range where the approximation is valid, the true nonlinear system might be explosive or tend toward another steady state, so the approximation would become permanently bad at that time.
- So with unbounded stochastic disturbances, our claims about the approximation take the form, "If σ is small enough, the probability that the linearized model is accurate over the time span $t = 0, \ldots, T$ is greater than $1 - \delta$." And by making σ small enough, we can make T as large and δ as small as we like.
- This kind of approximation claim can be made regardless of the eigenvalues of $D_1 f$.
- Nonetheless, the eigenvalues matter. If they are all small in absolute value, the T we can obtain with a given σ and δ will grow rapidly as $\sigma \to 0$. If the largest is much bigger than one, the T will grow extremely slowly as $\sigma \to 0$,

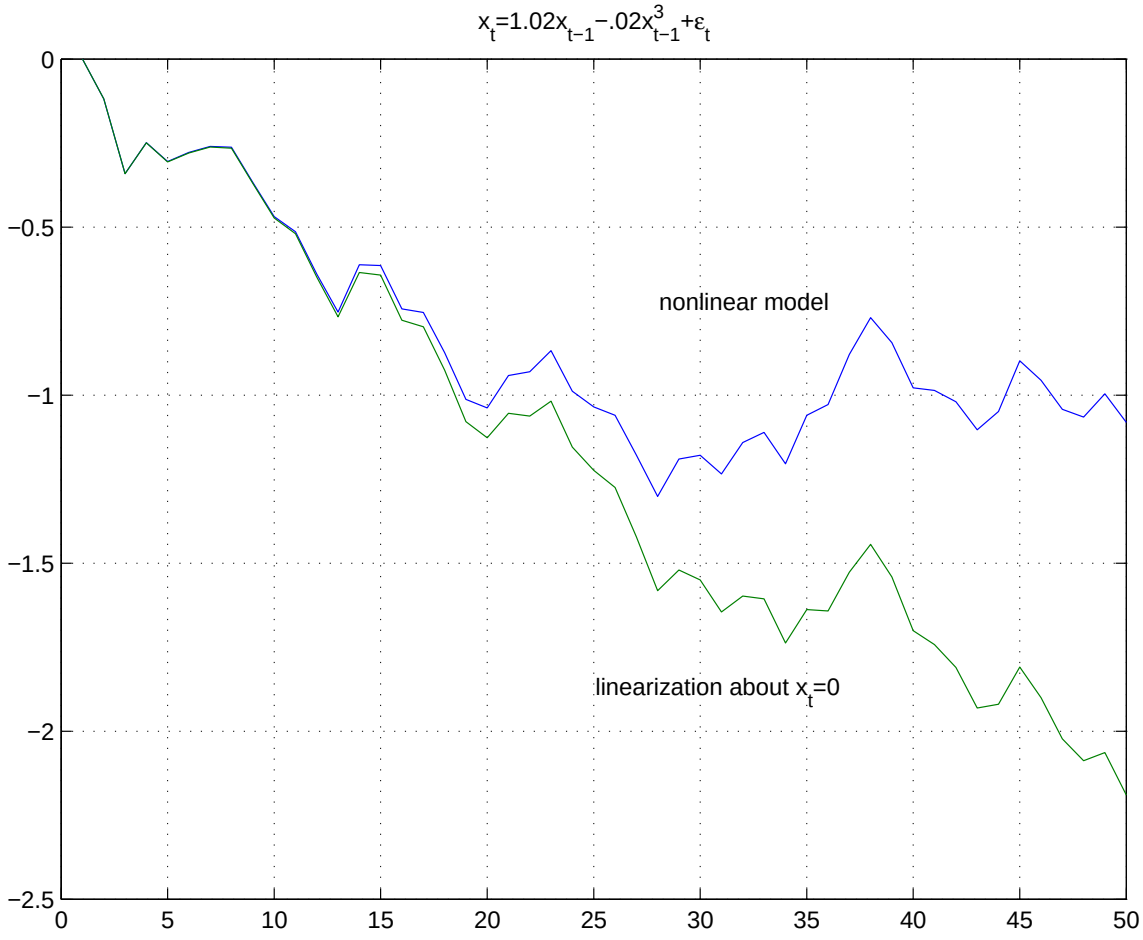

FIGURE 1

so that the approximation in practice is only good over very short time spans. Roots near one in absolute value are intermediate cases. They will allow T to grow as $\sigma \to 0$ much faster than for roots much greater than one, but much slower than for roots much less than one.

5. TWO EXAMPLES

Here are two examples of simulated time paths of nonlinear models and their linearizations. In one, the linearization is stable and the nonlinear model globally unstable, while in the other the linearization is unstable but the nonlinear model is globally stable. In both the linearization is quite accurate for a time (note the very different amounts of time), but then deteriorates catastrophically.

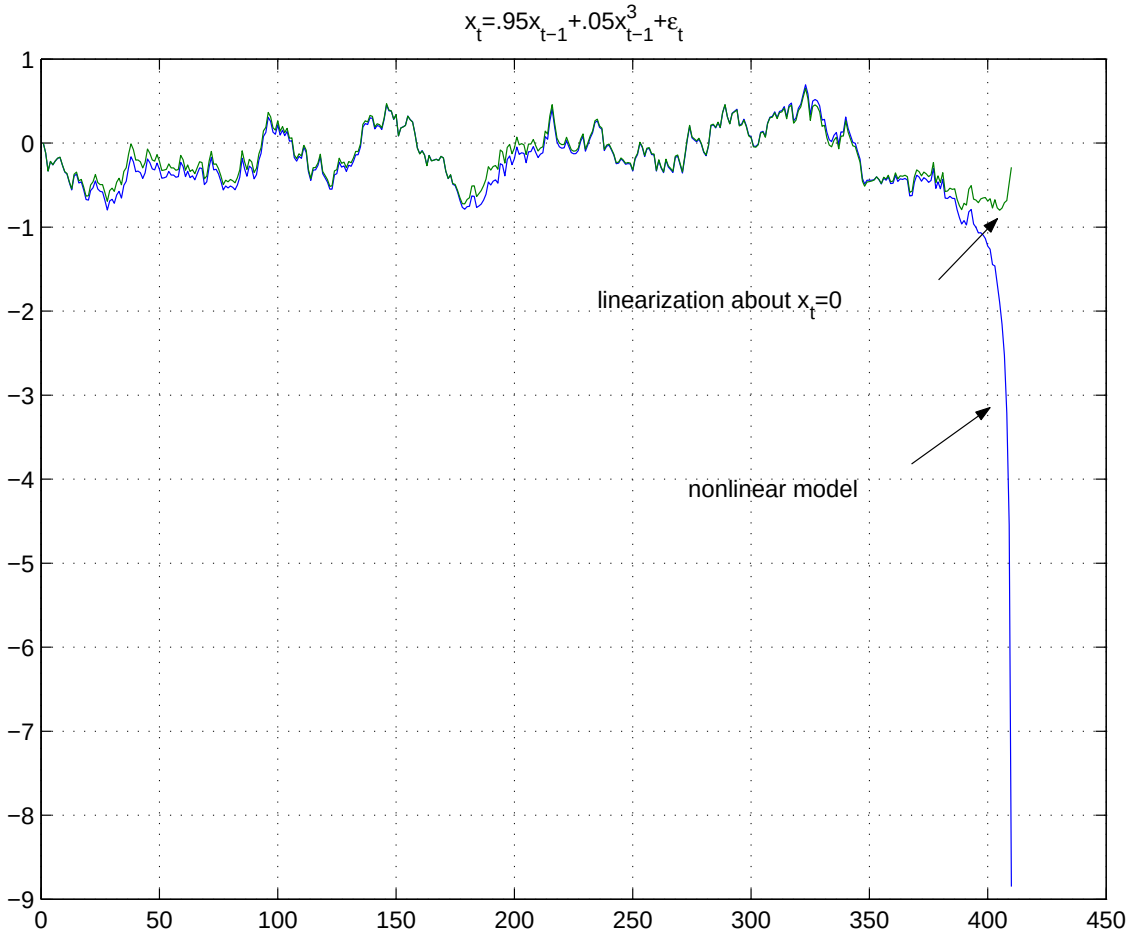

FIGURE 2

6

7. EXAMPLES: LINEARIZATION PITFALLS

This prescription seems simple, but it is not satisfied by some apparently reasonable procedures. For example, in a constrained optimization problem the prescription requires forming the system consisting of the first order conditions and the constraints and linearizing it. This is not the same thing as taking a second-order Taylor expansion of the objective function and a linear expansion of the constraints and solving the resulting linear-quadratic problem.

Example: A substitute LQ model: Suppose we want to solve

$$\max_{x,y} e^{-\frac{1}{2}(x^2+y^2)} \tag{5}$$

s.t.

$$(x-2-\varepsilon)^2+(y-2-\xi)^2 \leq 1\,. \tag{6}$$

The nonlinear equations defining a solution are the constraint (6) and the FOC's

$$-xe^{-\frac{1}{2}(x^2+y^2)} = (2x-4-2\varepsilon)\lambda \tag{7}$$

$$-ye^{-\frac{1}{2}(x^2+y^2)} = (2y-4-2\xi)\lambda\,. \tag{8}$$

We can eliminate λ by taking the ratio of the two FOC's to arrive at

$$\frac{x}{y} = \frac{2x-4-2\varepsilon}{2y-4-2\xi} \quad \Rightarrow \quad \frac{x}{y} = \frac{2-\varepsilon}{2-\xi}\,. \tag{9}$$

It is easy to understand the nature of the exact solution geometrically. Our objective function has circles centered at 0 for level curves, and the constraint is that the solution must lie on or inside a circle of radius 1 centered at $(2+\varepsilon, 2+\xi)$. The tangency of the constraint circle with a level curve always occurs on the line connecting the center of the constraint circle with the origin — this is exactly (9). We can see, then, that increasing ε, which increases the x-coordinate of the constraint circle center, will increase x by more than it increases y. It moves the solution farther from the origin, but also puts it on a ray from the origin with smaller slope (treating x as measured along the horizontal axis). The linear expansion of the equation system formed by (6) and (9), which gives a locally accurate answer, accordingly implies

$$dx = .82\varepsilon + .18\xi \tag{10}$$

$$dy = .18\varepsilon + .82\xi\,, \tag{11}$$

which fits our geometric description of the nature of the solution.

If we replace the objective function by a quadratic approximation to it, the level curves will still be circles centered at the origin, so there is no distortion in the problem from this source. But linearizing the constraint produces just

$$dx + dy = \varepsilon + \xi\,. \tag{12}$$

Obviously no matter how ξ and ε move about, this constraint remains a straight line with slope -1. Therefore its tangencies with level curves of the objective function all fall on the 45 degree line through the origin. The solution of this problem will imply that both ε and ξ have equal effects on dx and dy, which is far from being correct. Figure 3 displays the geometry, showing the optimum attained for $\xi = \varepsilon = 0$, the new optimum when $\xi = 0$, $\varepsilon = .2$, and the two constraint sets and two rays through the origin. Note that in this figure

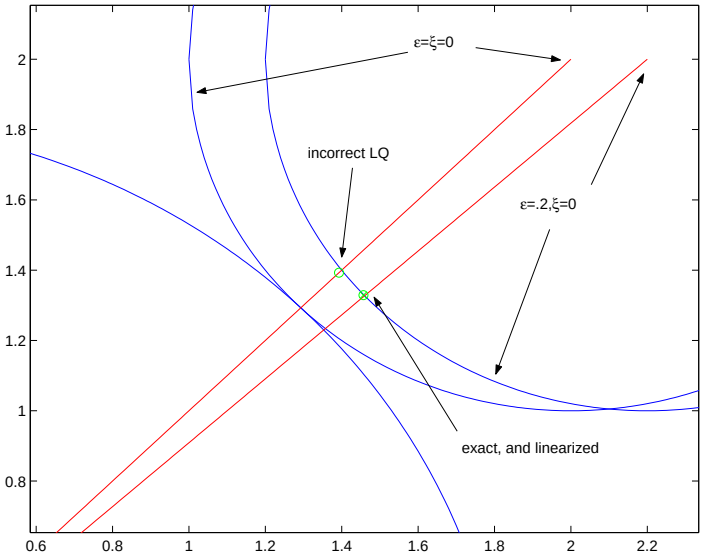

FIGURE 3. Geometry of error for the incorrect LQ solution

the linearized solution is so close to the true solution that the difference is not visible to the eye, while the incorrect LQ solution is visibly different.

It may seem odd that the correct linearization gives the right answer, despite the fact that it is a linear system that, like the incorrect LQ approximation discussed above, includes the linearized constraint (12). The difference is that the correct linearization involves, in the linearized FOC, second-order properties of the constraint. However, as we will see below, the fact that the correct linearization still uses the linearized constraint equation can create problems if we try to evaluate welfare properties of the solution to a stochastic model, because for this we usually need better than first-order accuracy in the solution.

8. CONDITIONS FOR ACCURATE WELFARE EVALUATION WITH LINEARIZED MODELS

Suppose we have a model solution in hand that gives us equations $x = h(\varepsilon)$ defining x as a function of ε. We would like to evaluate an objective function $E[U(x, \varepsilon)]$, perhaps because h represents one policy or market structure that we would like to compare to others. We are interested in whether we can get locally accurate results for small ε by replacing h with its first-order Taylor expansion about $\varepsilon = 0$.

The second-order Taylor expansion of U about $\varepsilon = 0$ is

$$U + \big(D_1 U\, D_\varepsilon h + D_2 U\big)\varepsilon + \frac{1}{2}\varepsilon' D_\varepsilon h' D_{11} U\, D_\varepsilon h \varepsilon + \overbrace{\frac{1}{2}\varepsilon' D_1 U\, D_{\varepsilon\varepsilon} h\, \varepsilon'}^{\dagger} + \varepsilon' D_{21} U\, D_\varepsilon h \varepsilon + \frac{1}{2}\varepsilon' D_{22} U\, \varepsilon\,. \quad (13)$$

Notice first that the first-order terms in the expansion will vanish when we take expectations, so that welfare comparisons must either be determined by the 0'th order term — meaning that they can be done with local accuracy by considering differences in deterministic solutions alone — or else they will depend on second or higher order terms in the expansion of U.

Next, note that the term singled out with the overbrace and marked † depends on $D_{\varepsilon\varepsilon} h$. Thus knowing an accurate first-order expansion of h is not enough, for second-order accuracy, unless $D_1 U(\bar{x}, 0) = 0$. There is of course one leading example of a case in which this would be true: that in which $x = h(0)$ is the solution to the equations $D_1 U(x, 0) = 0$, which is the FOC of an unconstrained maximization of $U(x, 0)$. Since many economic models can be formulated as maximization problems, this special case can be useful in practice.

However, economic models, if they are maximization problems at all, are most commonly *constrained* maximization problems. The † term does not drop out for the solution to constrained maximization problems. Why not, one might well ask, since constrained maximization problems can generally in principle be transformed into unconstrained problems by solving the constraints to produce a new, unconstrained problem with fewer variables. The answer is that the second-order properties of the constraint matter for the second-order expansion of U in ways not captured by the first-order solution. If the constraints are solved, they become embedded in the objective function and all their second-order properties are used. If the linearized solution for a longer list of variables is used, some of the second-order properties of the constraints are lost.

Note that it is nonetheless accurate to linearize the full set of FOC's and constraints, including in the system all the linearized constraints and all the Lagrange multipliers. Solving this system does yield a first-order accurate solution for all variables. The inaccuracy arises if this linearized solution for all variables is substituted into U in evaluating welfare, or if, equivalently, the linearized form of the constraints

is used to reduce the number of variables before substitution into U. To obtain accurate welfare evaluations, one must solve the original nonlinear form of the constraints, or a second-order approximation to them, to reduce the variable count before evaluating U.[1]

An example of this type of pitfall has been provided by Kim and Kim (2003). They consider the problem, which they interpret as arising from a simple model of international risk sharing,

$$\max_{c_1,c_2}\{c_1 + c_2\} \tag{14}$$

s.t.

$$e^{c_1} + e^{c_2} = e^{y_1} + e^{y_2}\,. \tag{15}$$

The FOC's for this problem[2] are

$$1 = \lambda e^{c_1} \tag{16}$$

$$1 = \lambda e^{c_2} \tag{17}$$

which lead to the intuitively obvious solution

$$e^{c_1} = e^{c_2} = \frac{e^{y_1} + e^{y_2}}{2}\,. \tag{18}$$

The linearized version of the constraint and FOC's, expanded about $y_1 = y_2 = c_1 = c_2 = 0$, is

$$dc_1 + dc_2 = y_1 + y_2 \tag{19}$$

$$d\lambda + dc_1 = 0 \tag{20}$$

$$d\lambda + dc_2 = 0\,. \tag{21}$$

which leads to $dc_1 = dc_2 = (y_1 + y_2)/2$.

[1]The idea of using the linearized form of the model to solve it, but then simulating it or otherwise evaluating expected welfare using the original nonlinear constraints, has been suggested and called called "partial linearization" by Evans and Marshall (1998), though they do not justify it as giving a second order approximation to welfare. Indeed they apply it in a model of term structure, where the fact that the solution for the variables themselves is still only first-order accurate, and thus certainty-equivalent, may be a serious problem.

[2]In this form, the objective function is not concave, or even quasi-concave, for $0 < \gamma < 1$. However, the problem is a transformation of one written in terms of $C_i = e^{c_i}$ that does have concave objective function and convex constraint for any $\gamma > 0$, so we can be confident that the FOC's define an optimum.

Thus so long as $E[y_1] = E[y_2] = 0$, the linearized solution implies that expected welfare is zero, and independent of the variance of y_1 and y_2. But if we take a second-order expansion of the welfare function using the true solution (18), we arrive at

$$y_1 + y_2 + \frac{(y_1 - y_2)^2}{4}\,. \tag{22}$$

Clearly then welfare does in fact depend — positively — on the variance of y_1 and y_2.[3]

The solution we suggested above, solving the constraint to eliminate one of the variables, works here. If we use the linearized solution $dc_1 = (y_1 + y_2)/2$ for c_1, then substitution into the exact constraint (15) produces

$$dc_2 = \log(e^{y_1} + e^{y_2} - e^{(y_1+y_2)/2})\,, \tag{23}$$

which has the second-order Taylor expansion

$$dc_2 = \frac{y_1 + y_2}{2} + \frac{(y_1 - y_2)^2}{4}\,. \tag{24}$$

It is then clear that the second order expansion of the objective function $c_1 + c_2$ is exactly the true expansion given in (22).

Note that this procedure has not resulted in a second-order accurate solution for the individual variables c_1 and c_2. It has only produced a first-order accurate solution which, substituted into the welfare function, produces a second-order accurate approximation to welfare.

9. Integral Constraints

The problems we have discussed to this point, though we have applied the expectation operator to their welfare functions, are based on deterministic equation systems: each exogenous ε value produces a solution for x which can be found by solving a set of nonlinear equations for a real vector x, given the real vector ε, one ε, x pair at a time. More common in economics are problems in which some constraints involve expectations, so that they can only be evaluated if the entire mapping from the space of ε's to the space of x's is known. Such problems can also be handled by linearization methods, but they raise some new issues.

Perhaps most importantly, in these problems, because the constraints and FOC's do not simply form a nonlinear equation system in real variables, the constraints cannot be solved to convert the problem to unconstrained form. To take the simplest possible example, consider the problem

$$\max_x E[-\tfrac{1}{2}(x - 1 - \varepsilon)^2] \tag{25}$$

[3]The fact that increased variance increases welfare comes from the fact that the expectation of the right-hand-side of the constraint, $E[e^{y_1} + e^{y_2}]$, is increasing in the variance of y_1 and y_2.

subject to

$$E[e^x] \leq 1\,. \tag{26}$$

The first order condition for the problem is

$$-x + \varepsilon + 1 = \lambda E[e^x]\,, \tag{27}$$

where it is important that λ is a constant, not varying with ε. Assuming the constraint is binding and using (26) in (27) gives us

$$E[e^{1+\varepsilon-\lambda}] = 1\,. \tag{28}$$

A second order expansion of this equation, or an assumption that ε is normally distributed, gives us

$$\lambda \doteq 1 + \frac{\sigma^2}{2} \tag{29}$$

and therefore the conclusion

$$x \doteq \varepsilon - \frac{\sigma^2}{2}\,. \tag{30}$$

This exact (for the Gaussian ε case) or second order solution illustrates an important general property of solutions to models with integral constraints. The first-order solution (here simply $x = \varepsilon$) differs from the second-order solution (here (30)) not just in a "second order" term involving cross products of ε's (which actually does not occur at all here), but in the "0'th-order" term, the constant. This can lead to some semantic confusion over what is meant by a "first-order accurate" solution. The natural definition, though, looks for accuracy as we shrink, together, both the size of the particular ε for which we solve in (26) and (27) and the spread of the distribution of the ε over which we integrate in taking the expectations. Thus the term in σ^2 in (30) is of second-order, and the first-order solution, $dx = \varepsilon$, is indeed first-order accurate.

However, as usual, first-order accurate solutions do not provide second-order accurate welfare assessments, and thus are likely often not to be useful at all for welfare assessment. Here the second-order solution tells us that expected utility is approximately $-\frac{1}{2}(1 + \sigma^2/2)$, while the first-order solution tells us that expected welfare is the same as for the deterministic $\varepsilon \equiv 0$ case.

Unlike the case without integral constraints, however, this model does not provide us a straightforward way to make a first-order solution second-order accurate. We cannot solve the constraint to reduce the number of variables, because the constraint is not an ordinary equation, but a functional equation.

To justify use in welfare comparisons of linear approximations to solutions of systems involving expectational constraints, we have to assume that the constraints themselves are of small importance. For example, in our current example we could replace the constraint $E[e^x] \leq 1$ with $E[e^x] \leq e^{1-\gamma}$, with γ itself assumed to be first-order small. Since at $\gamma = 0$ the constraint becomes non-binding, assuming γ small

makes λ nearly zero and makes inserting the first-order solution into the objective (25) produce a second-order accurate approximation to it. Proceeding this way obviously evades mathematical difficulties at a potentially serious substantive cost. For example, Woodford (2002) in taking this approach is led at one point to impose the assumption that the economy he studies is near the satiation point in real balances (which requires either deflation at near the real interest rate or interest payments on cash and high-powered money at near the real rate of interest) or a negligible role for real balances in the economy. As Woodford acknowledges, these assumptions imply important limitations on the range of applicability of the results obtained by the first-order methods.

10. DYNAMIC MODELS

Most macroeconomic models are dynamic, so that they do not fit directly into any of the simple setups we have discussed so far. To cast them into the form of a standard equation-solving problem, we give the equation system this form:

$$F(x_t, x_{t-1}, \varepsilon_t) = 0 \tag{31}$$

$$E_t(G(x_{t+1}, x_t, \xi_t) = 0\,. \tag{32}$$

Assuming the model has a deterministic steady-state solution $\bar{x}$, i.e. that

$$F(\bar{x}, \bar{x}, 0) = 0 \tag{33}$$

$$G(\bar{x}, \bar{x}, 0) = 0\,, \tag{34}$$

we can look for a linearization around this steady state for small ε and ξ. A solution will have the form $x_t = h(x_{t-1}, \varepsilon_t, \sigma^2)$. We look for a second-order solution by substituting $h(x_t, \varepsilon_{t+1}, \sigma^2)$ for x_{t+1} in (32) and $h(x_{t-1}, \varepsilon_t, \sigma^2)$ for x_t in (31). If we expand the resulting system in a second-order Taylor expansion in x, ε, and σ, we obtain a system in which we can hope to determine the first and second order coefficients in the expansion of h by equating to zero the first and second order coefficients in the expansion of (31) and (32).

The details of doing this are a mess even for simple dynamic models. Nonetheless the method is in principle straightforward, so where it is important to obtain second-order accuracy, the method is feasible. An example of application to a simple growth model is in Judd (1998, section 13.7). There are as yet few (no?) examples in the literature of a full-fledged

REFERENCES

EVANS, C. L., AND D. A. MARSHALL (1998): "Monetary Policy and the Term Structure of Nominal Interest Rates: Evidence and Theory," *Carnegie-Rochester Conference Series on Public Policy*, 49, 53–111.

JUDD, K. L. (1998): *Numerical Methods in Economics*. MIT Press, Cambridge, Mass.

KIM, J., AND S. KIM (2003): "Spurious Welfare Reversals in International Business Cycle Models," *Journal of International Economics*, 60, 471–500.

WOODFORD, M. (2002): "Inflation Stabilization and Welfare," *Contributions to Macroeconomics*, 2(1), Article 1.