

Discussion of ‘Natural Expectations, Macro Dynamics, and Asset Pricing’.

Martin Eichenbaum*,

August 2011[†]

1 Introduction

This last July marked the fiftieth birthday of John Muth’s seminal paper, ‘Rational Expectations and the Theory of Price Movements.’ This assumption about how expectations are formed is now the standard approach in modern, quantitative macro. But there is an increasingly important literature which examines the deeper foundations of the rational expectations hypothesis and the assumption that agents have common priors.¹ See for example Acemoglu and Ozdaglar (2010) for a review of recent research on belief and opinion dynamics in social networks. If we had an empirically convincing model of asset pricing and financial markets that embodied simple forms of rational expectations, I suspect that this literature would be of interest primarily to a small set of decision theorists. As recent events have forcefully brought home, we do not. So we are forced to re-assess all of our cherished maintained assumptions. Like Loehman Brothers, the luxury of sacred cows is dead.

Fuster et al. take it as given that it not reasonable to assume that agents have common priors and rational expectations when we try to explain asset prices. They then investigate how the economy would behave if (1) ‘fundamentals’ are hump shaped, exhibiting momentum in the short run and partial reversion in the long run, and (2) agents forecast fundamentals using ‘reasonable models that are too simple to capture long-term dynamics of key time series.’ These forecasts are based on simple univariate time series models and correspond to what the authors call ‘natural expectations.’ The associated specification error leads agents to systematically over-estimate the persistence of shocks. Fuster et. al. argue that the natural expectations assumption allows them to ‘generate examples that are consistent with a host of asset pricing phenomenon,’ such as the equity premium puzzle and the apparent volatility of asset prices relative to fundamentals.

My discussion of the paper is organized as follows. First, I review the authors’ statistical argument in favor of ‘natural expectations.’ This argument revolves around the same feature

*Northwestern University, NBER, and Federal Reserve Bank of Chicago.

[†]I thank Ben Johannsen for superb research assistance in preparing this comment.

¹See for example Acemoglu and Ozdaglar (2010) for a review of recent research on belief and opinion dynamics in social networks.

of the data emphasized in the literature that explores the problems of agents who are vitally concerned with the possibility of model misspecification.² This feature is that there is an enormous amount of statistical uncertainty about the true data generating process of objects which agents care about. In principle, this observation is very damaging to the Fuster et al. paper. Why would risk averse agents ever settle on a particular forecasting model which is very likely to be misspecified and act as if the model is, with, probably one, true?

To explore the seriousness of this problem, I investigate how much agents would pay to know the true model for ‘fundamentals.’ To my surprise, I find that the utility cost of specification errors associated with natural expectations can be quite small. This result holds even though the implications of the specification error association with natural expectations for key moments of the data can be quite large. I make this argument using a variant of the simple setup in John Cochrane’s classic paper, ‘The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near Rational Alternatives’ (1989).

Do my misspecification ‘cost’ numbers holds up in the actual Fuster et al. model? I don’t know. Absent such a calculation, the authors’ statistical results about the difficulty of pinning down the ‘true’ ARIMA representation of fundamentals can be thought of as evidence *against* the way they model agents’ behavior. In fact, for reasons discussed below, I suspect that these costs will be small in their application. The reason might be anticipated from Cochrane (1989) and Lucas’ (1987) famous argument that the welfare costs of business cycles are small. In Cochrane (1989), Lucas (1987) and Fuster et al. model agents respect their intertemporal budget constraints. So in a well defined sense, their average level of consumption would not be different if we eliminate business cycles (Lucas), or agents used optimal decision rules (Cochrane) or agents had rational expectations (Fuster et al.). But as in Cochrane (1989), the change in the decision rule induced by the assumption of natural expectations can have large implications for various moments of the data. If my conjecture is correct, it would provide the basis of a very powerful defense of the approach used in this paper to resolve asset pricing puzzles.

2 Natural Expectations

According to the authors, ‘The premise of our approach ... is that economic agents tend to make forecasts based on statistical or mental models that are reasonable given the data available to them, but “too simple” to fully capture the long-term dynamics of many economic time series.’ These simple model leads agents to over-estimate the persistence of shocks when the true process is hump shaped.

To consider the statistical issues involved, suppose you want to forecast a time series, Δd_t . Also suppose that for unspecified reasons you consider only *univariate* models, say ARMA(p, q) specifications of the form

$$\alpha(L)\Delta d_t = \beta(L)\varepsilon_t.$$

Here $\alpha(L)$ and $\beta(L)$ are polynomial operators in the lag operator of order p and q , while ε_t is uncorrelated with lagged values of Δd_t or ε_{t-s} , $s \geq 1$. The authors are interested

²For recent reviews of this literature see Hansen and Sargent (2010) and Epstein and Schneider (2010).

in Δd_t processes that exhibit hump shaped dynamics, i.e. processes that exhibit positive autocorrelation in the short run, but partially revert to the mean in the long run. Their measure of the persistence in Δd_t is

$$\Gamma = \frac{1 + \beta(1)}{1 - \alpha(1)}.$$

Recall that the spectral density of Δd_t at frequency zero is

$$S_w(0) = \Gamma^2 \sigma_\varepsilon^2.$$

This object is notoriously difficult to estimate with precision in small samples (see for example Christiano, Eichenbaum and Vigfusson (2006)). It follows that it should be hard to estimate Γ with much precision.

Suppose we take as given that agents care about Γ per se. The statistical motivation in the paper for natural expectations focuses on the difficulty of getting p and q right based on a particular statistical metric like the AIC or BIC criterion. These are criteria for choosing models that are good for one-step ahead forecasting exercises. So it is not surprising that models based on such criteria could be very bad for long-horizon forecasts or estimating Γ . But if the latter play an important role in agents' decision problems, why would they use AIC-like criteria?

More importantly, in the model, agents must forecast Δd_{t+s} , $s \geq 1$, to solve an *economic* problem as opposed to a purely statistical one. In general there is no reason to suppose they would separate their forecasting and control problems. Fuster et al. just take this separation property as given in the sense that agents ignore both model and parameter uncertainty. By assumption, agents in the model first choose an ARMA(p , q) model for Δd_t . Next, they estimate the parameters of their preferred statistical model. Finally, they solve their economic optimization problem, taking the ARIMA specification and point estimates as true. This procedure seems very hard to defend. In general we would expect agents to take both model and parameter uncertainty into account when making their plans. Indeed, the more uncertainty there is about the data generating process, the more skeptical we ought to be about the authors' separation assumption.

3 Statistical uncertainty

The measure of d_t which Fuster et al. use in their asset pricing application is the log of real capital income.³ By assumption the true data generating process for d_t is an ARIMA(40, 1, 0) which they estimate using data over the period 1947Q1-2010Q4. The implied point estimate of Γ for Δd_t is 0.43. Fuster et al. JEP (2010) conduct a similar exercise using data on the log of per capita real GDP which they assume is an ARIMA(0, 1, 12). The implied point estimate of Γ for Δd_t is 1.3. For both measures of d_t , Fuster et. al. argue that if agents use a parsimonious ARIMA process to fit the data, they will over-estimate the persistence of the process.

³See footnote 30 of the paper for data definitions.

A key feature of the data is that there is a great deal of statistical uncertainty about the true value of Γ . To quantify this uncertainty I take as given the authors' estimated ARIMA processes for the two measures of d_t , and generate 10,000 artificial time series, each of length 250. I then estimate the relevant ARIMA process for d_t using each time series and compute quantiles for the persistence statistic. The following table summarizes my results.

Real Per Capita GDP		Real Capital Income	
Quantile	Γ	Quantile	Γ
2.5%	0.50	2.5%	0.25
50%	1.3	50%	0.43
97.5%	2.04	97.5%	0.73

The key result here is that even if agents knew the true ARIMA for d_t , there would be a great deal of uncertainty about Γ .

Of course agents do not know the true ARMA(p, q) process for Δd_t . The authors stress that agents might use 'simple' univariate tests like AIC to determine p and q . But if agents proceeded this way, the most important thing that they would learn from the data is that there is enormous amount of uncertainty about the values of p, q and Γ . To make this point concrete, I suppose that the true data generating process for d_t is the estimated ARIMA(40, 1, 0) for the log of capital income. I then simulated 100 synthetic time series of size 250. For each synthetic time series, I estimated an ARIMA($j, 1, 0$), $j = 1, \dots, 40$, and calculated the preferred representation using the AIC. Figure 1 summarizes the cumulative distribution function for the preferred specification. It is evident that from the perspective of agent using the AIC criterion, there is a great deal of uncertainty about the true value of p . Of course this uncertainty maps into enormous uncertainty about the response of Δd_t to a shock.

Viewed as a whole the previous results make clear that the agents living in the authors' model face enormous uncertainty about features of the data that would seem to be critical to them as decision makers. In the face of such uncertainty why would we assume that agents condition their decisions on a particular statistical model?

4 The key mechanism in the paper

Before providing a possible defense for the authors' assumptions it's useful to provide some intuition for the key mechanism in their model. The authors consider a small open endowment economy with two assets. The first asset is foreign debt, b_t , which can be held at a constant interest rate R ,

$$\beta R = 1.$$

Here β is agents' common discount rate. The second asset is a domestic equity tree, which generates a dividend, d_t , governed by

$$\Delta d_t = \rho_1 \Delta d_{t-1} + \dots + \rho_n \Delta d_{t-n} + \varepsilon_{dt},$$

where n is a non-negative integer. The equity tree must be held by a representative domestic agent who has an exponential utility function with habit formation governed by the parameter $\gamma \in (0, 1)$.

Since the domestic agent can borrow and lend at a constant interest rate, we have the first order condition

$$u'(c_t, c_{t-1}) = E_t u'(c_{t+1}, c_t).$$

If $\gamma = 0$, we obtain

$$u'(c_t) = E_t u'(c_{t+1}).$$

Since agents do not have rational expectations, $E_t(\cdot)$ does not correspond to the conditional expectations operator defined in terms of the true data generating process. So with natural expectations, the model does not predict martingale-like behavior for consumption even when $\gamma = 0$.

In this environment consumption is a weighted average of lagged consumption and the (risk-neutral) annuity value of perceived future dividends, x_t , down-shifted by a constant φ that involves a correction for risk aversion.

$$c_t = \frac{\gamma}{R} c_{t-1} + \left(1 - \frac{\gamma}{R}\right) x_t - \varphi.$$

The equilibrium price of the Lucas tree is given by,

$$p_t = \sum_{s=0}^{\infty} \frac{E_t d_{t+s}}{R^s} - \frac{R\alpha \text{Var}_t(\Delta c_{t+1})}{\left(1 - \frac{\gamma}{R}\right)(R-1)^2}.$$

Note that when $\gamma = 0$, agents *plan* to respond to an innovation in d_t by raising consumption on a one time basis. The larger is the present value of the change in d_t , the larger will be the change in consumption and the larger will be the response of p_t .

Taking as given the true data generating process, the order of p that agents use when estimating the law of motion for Δd_t is critical in determining inferences about persistence, the change in the present value of d_t induced by a shock to Δd_t , the change in consumption and $\text{Var}_t(\Delta c_{t+1})$. The lower is p , the higher is agents' point estimate of Γ . So the larger will be the initial change in consumption and the perceived riskiness of the Lucas tree.

To understand the last result, suppose that agents act as if Δd_t is governed by the AR(40, 0) model that Fuster et al. estimate from the data. Suppose that we simulate a time series of length 10,000 and estimate an ARMA(p, 0), $p < 0$ using the simulated data. In this way we calculate an approximation to the probability limit of the coefficients of the misspecified ARMA(p, 0) model. Figure 2 displays the impulse response function for d_t implied by different values of p . Note that persistence and the implied expected change in the present value of d_t induced by a shock is *decreasing* in p .

Not surprisingly, when agents use values of $p < 40$, they increase consumption by more than they would if they used the true value of p . The actual present value of the response of consumption must be the same across different values of p because agents' intertemporal budget constraint must hold. The way this happens is that over time, working through the lens of their incorrect statistical model, agents infer incorrectly that there is a sequence of negative shocks to earnings. So they decrease consumption relative to its peak value. It follows that even when $\gamma = 1$ and $\beta R = 1$, consumption will not be a martingale. More generally, a classic Hansen and Singleton (1982) type analysis would correctly reject the representative household rational expectations model of asset pricing. The underlying failure

of the model isn't heterogeneity among agents or financial frictions. It is simply the failure of the rational expectations assumption.

Hopefully the previous discussion makes clear the intuition for why a 'wrong' value for p affects the volatility of consumption. The same logic makes clear that agents ought to worry about model uncertainty, i.e. they would pay a positive amount to know the true ARMA model for Δd_t . If this amount is large, I do not see how the authors can justify abstracting from the issue of how agents would act if they acknowledged model and parameter uncertainty. Absent some justification for the separation assumption, it is hard to have much confidence in the asset pricing implications of the model.

5 Cochrane (1989) to the rescue?

Cochrane (1989) shows that for an interesting class of problems, feasible first order mistakes in the choice variables of a single agent problem have second order utility consequences. Moreover these deviations can have a dramatic impact on the time series implications of consumption allocation models and the behavior of asset prices. Is that what's going in this paper? If so, we have a possible rationale for the authors' assumption that agents don't worry about model misspecification. Indeed the paper could be viewed as an interesting example of the Cochrane principle with the extra, important refinement that we can draw on the psychology literature to motivate the particular 'small' mistake that agents are making.

To illustrate this point, consider the following modified version of one of Cochrane's examples. Consider a representative household with exogenous income given by

$$(1 - \rho L)d_t = \bar{y} + \rho_1(1 - \rho L)d_{t-1} + \dots + (1 - \rho L)\rho_n d_{t-n} + \sigma \varepsilon_t.$$

The parameter ρ can be arbitrarily close to 1 so that we that we can mimic the times series behavior of ARIMA(p, 1, 0) processes well. The household maximizes the criterion

$$U = \frac{-1}{2} \sum_{t=0}^{\infty} \beta^t (c_t - \bar{c})^2,$$

where \bar{c} is a nonnegative constant. The household can borrow and lend at the constant interest rate

$$R = 1 + r, \quad \beta = \frac{1}{1 + r}$$

where r is a nonnegative constant. The household's budget constraint is given by

$$k_{t+1} = (1 + r)k_t + d_t - c_t.$$

Here k_t is beginning of period t accumulated capital or nonhuman wealth, and plays the role of bonds in Fuster et. al..

Cochrane (1989) considers the case where $\rho_1 = \rho_2 = \dots \rho_n = 0$. In this case, d_t is an AR(1) and optimal consumption evolves as

$$\begin{aligned} c_t^* &= rk_t^* + m^*(d_t - Ed) \\ k_{t+1}^* &= (1 - m^*)(d_t - Ed) \\ m^* &= \frac{r}{1 + r - \rho}. \end{aligned}$$

Here Ed denotes the unconditional expectation of d .

Motivated by Flavin’s (1980) findings, Cochrane supposes that the agent’s decision rule for consumption is given by

$$\begin{aligned} c_t^+ &= rk_t^+ + m^+(d_t - Ed) \\ k_{t+1}^+ &= (1 - m^+)(d_t - Ed) \\ m^+ &> m^*. \end{aligned}$$

By construction, the alternative decision rules, incorporating different values of m^+ , respect the agent’s budget constraint. Cochrane chooses m^+ to account for Flavin’s (1980) influential estimate of the excess-sensitivity of consumption to current income. By construction, the model ‘solves’ the excess-sensitivity puzzle.

Cochrane emphasizes that this solution is only interesting if the required deviation from the benchmark value of m^* is small in the sense that an agent wouldn’t pay much to go from m^+ to m^* . In fact Cochrane shows that the agent pay only a trivial amount (substantially less than \$1 a year) to go from m^+ to m^* . He argues that a similar logic applies to various asset pricing puzzles.

We can perform a similar experiment to evaluate the deviation from rational expectations adopted by Fuster et al. Suppose that the law of motion for $(1 - \rho L)d_t$ is the $AR(40)$ corresponding to the specification for Δd_t used in the paper. Also assume that agents incorrectly think that $(1 - \rho L)d_t$ is governed by a low order AR processes. In addition I use the Cochrane calibration which makes relative risk aversion equal to 4 in an area around steady state. This value is the same as the one used by Fuster et. al.

Panel A of Figure 3 displays the impulse response function of consumption for various value of p . For reference, Panel B of Figure 3 shows the impulse response function of consumption if agents assume $p = 4$, but I simply impose different value of m^* as in Cochrane.⁴ The key results from Panel A are as follows. First, under rational expectations, consumption adjusts once and for all and that’s the end of the story. The rise in consumption induced by a unit shock to income is about 1/3. Second, consistent with Fuster et. al., the more parsimonious is the AR process used by agents, the larger is the initial jump in consumption. For example in the $AR(2)$ case, the ratio of the rise in consumption to the innovation in earnings is about 1.2. As in Fuster et. al., when $p < 40$ agents are using a misspecified model and are over-estimating the present value of the change in income. So consumption has to drop from its peak level. Consumption asymptotes to a level lower than the long-run response if agents had been using the correct AR specification. This decline must occur to ensure that agents’ budget constraints hold.

The key question is: how much would agents pay to know the truth? The answer is between 35 cents and \$1.85 a quarter, depending on the AR that they mistakenly use. That’s pretty small. Of course lots of frictions could generate the kind of over-shooting in consumption discussed above, including the psychological factors emphasized by Fuster et al. Panel B, which allows for different values of m^+ generates similar response patterns, with higher values m^+ corresponding roughly to lower value of p . Presumably different values of m^+ could be generated using various credit market imperfections.

⁴I solve this version of the model by multiplying the coefficient on current d_t by a factor greater than its optimal value, one.


6 Conclusion

This a very interesting paper which got me to think about lots of fascinating issues. It is clear that the authors' perturbation of the standard model helps them account for certain asset pricing puzzles. The key omission of the paper is a defense of the separation procedure that they invoke when modeling agents' behavior. That said, my calculations suggest that the authors might be able to come up with a potentially compelling defense of this procedure. Agents may just not be willing to pay very much to know the true data generating mechanisms in their setup. My guess is that what is going on is that despite the fact that they are mistaken about the exact nature of the data generating mechanism, agents' consumption levels are, on average correct. By this I mean that their intertemporal budget constraint holds. To the extent that this defense of the separation procedure holds, it would one more application of Lucas' (1987) point that once you get the levels right, the fluctuations do not seem to matter all that much.

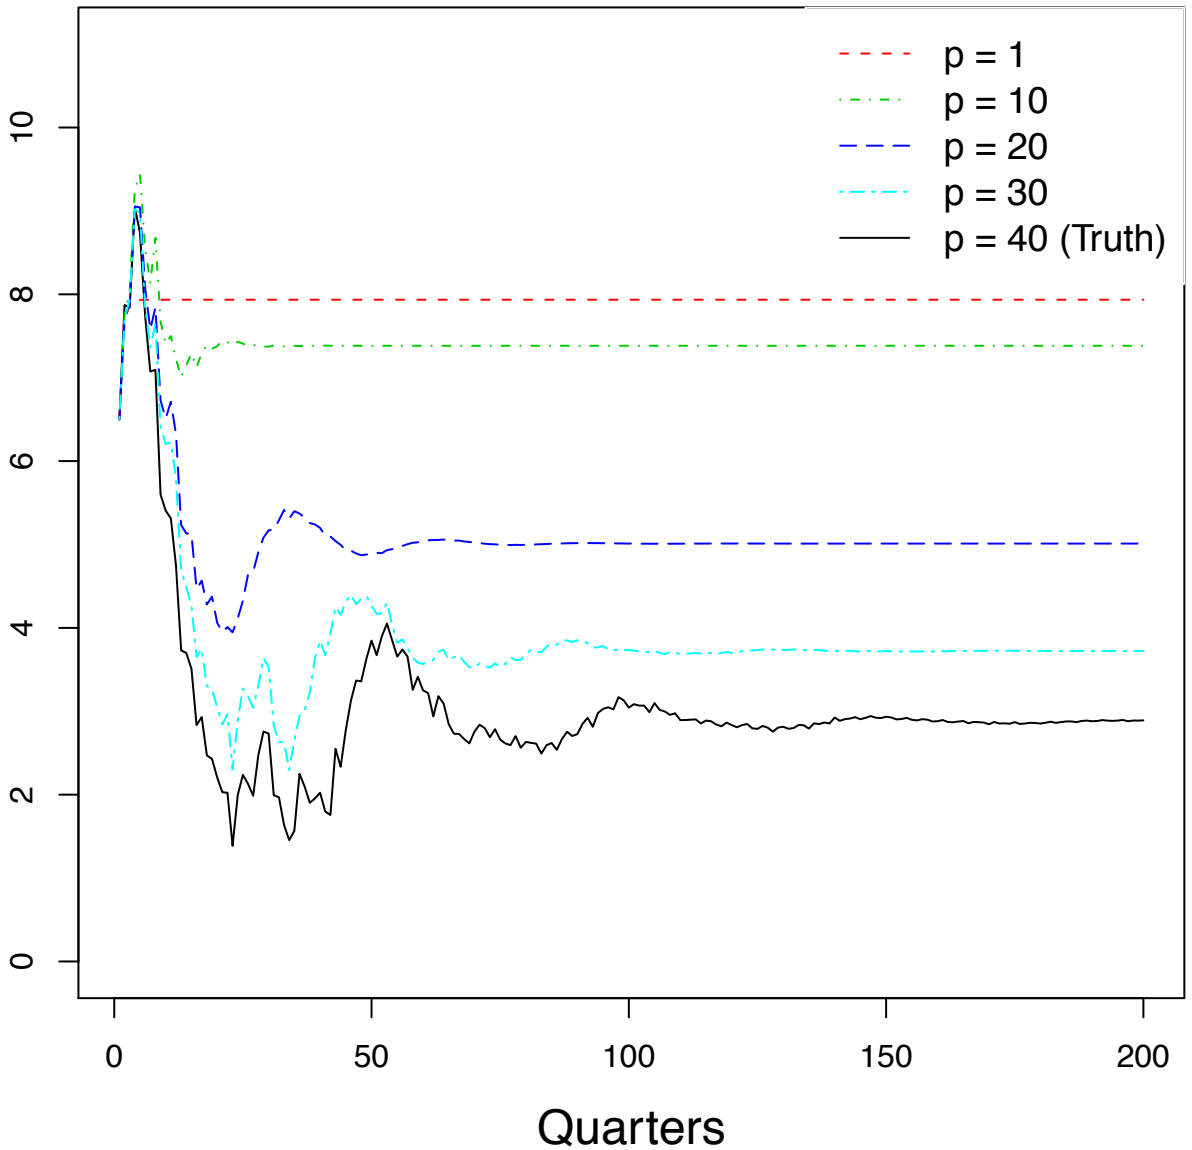
References

- [1] Acemoglu, Daron, and Asuman Ozdaglar, 'Opinion Dynamics and Learning in Social Networks', forthcoming, *Dynamic Games and Applications*,(2011)
- [2] Christiano, Lawrence, Eichenbaum, Martin and Robert Vigfusson, 'Assessing Structural VARs', *NBER Macroeconomics Annual 2006*, Volume 21, Daron Acemoglu, Kenneth Rogoff and Michael Woodford, editors.
- [3] Cochrane, John, "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives" *American Economic Review* 79 (June 1989) 319-337.
- [4] Epstein, Larry and Martin, Schneider, 'Ambiguity and asset markets', *Annual Review of Financial Economics*, 2010, vol 2.
- [5] Fuster, Andrea, Laibson, David and Brock Mendel, 'Natural Expectations and Economic Fluctuations', *Journal of Economic Perspectives*, Volume 24, Number 4, Fall 2010, pages 67-84.
- [6] Hansen, L.P. and Sargent, T.J, 'Wanting robustness in macroeconomics', *Handbook of Monetary Economics*, 2010, vol 3.
- [7] Lucas Jr., Robert E., 1987. *Models of Business Cycles*, Basil Blackwell, New York.

Cumulative Distribution Function of AIC Preference

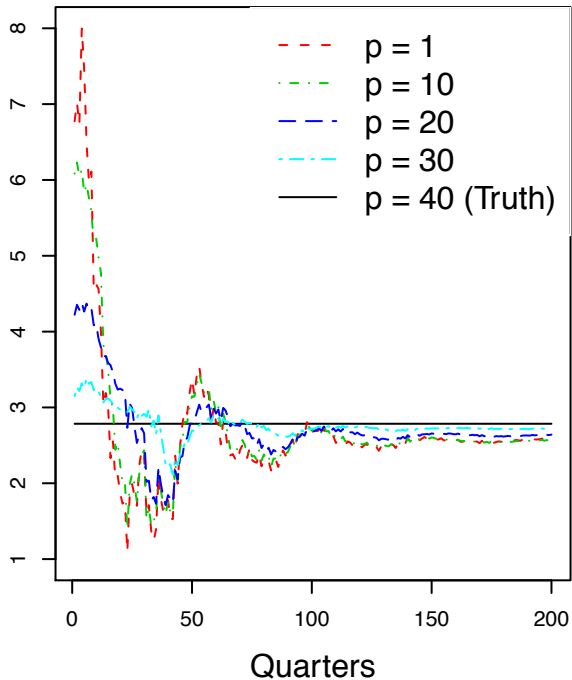


Response of Income In Levels



Impulse Response of Consumption

Improperly Specified ARIMA(p, 1, 0)



Excess Sensitivity to Current Income

