Time-to-Plan and Investment Dynamics

James A. Kahn and Louis Maccini[1]

Revised: December 30, 2015 (preliminary and incomplete)

[1]Kahn: Yeshiva University; Maccini: Johns Hopkins University

Abstract

A large body of research has established that investment responds sluggishly toward a long-term target. Quantitative models have contended with this by assuming various forms of adjustment costs, both in capital and in the rate of investment. We provide an alternative approach to modeling investment dynamics that focuses on the supply side, in which producers of investment goods use delivery lags and unfilled orders to smooth production and thereby modulate investment flows.

A large body of research has established that investment responds sluggishly toward a long-term target. Quantitative models of the demand for investment have contended with this by assuming various forms of adjustment costs, both in capital and in the rate of investment. While costs of adjusting capital are motivated by observed phenomena such as installation costs and disruption of production, costs of adjusting the rate of investment have been added to models more as an ad hoc device to generate hump-shaped impulse responses. See, for example, Basu and Kimball (2003), Chari, Kehoe, and McGratton (2000), and Christiano, Eichenbaum, and Evans (2005), as well as the substantial body of literature on the q theory of investment, such as Hayashi (1982).

We provide an alternative approach to modeling investment dynamics that focuses on the supply side of the market or the production of capital goods, based on the flexible time-to-build framework of Lubik, Sarte, and Schwartzman (2015, hereafter LSS). In this case, producers of capital goods face lags either in the time to build capital goods or time to plan for their production. Several papers build on the seminal work of Kydland and Prescott (1982) to focus on the time to build capital goods. These include the studies by Casares (2006), Edge (2007), Lucca (2007), Del Boca, Galeotti, Himmelberg and Rota (2008), and Tsoukalas (2011). Other studies focus on producers of capital goods that use use delivery lags and unfilled orders to smooth production and thereby modulate investment flows. This work builds on the work of Maccini (1973) and Carlton (1983) and includes more recent work by Reagan and Sheehan (1985), Haltiwanger and Maccini (1989), West (1989), Kahn (2010), and Nalewaik and Pinto (2015)

The next section of the paper presents the model that we develop and reports simulations of the model. A succeeding section examines data on shipments, new orders, and unfilled orders and compares their behavior with implications of the model.

1 Model with Flexible Order Lead Times

Suppose investment at date t, i_t, requires $S+1$ inputs from dates $t-S$ through t. Let $z_{t,t-s}$ denote the quantity of the input to i_t at date $t-s$, for $s=0,...,S$. For now we will not model the production of these inputs, but assume that the final good producer can purchase these at a given set of prices $p_{t,t-s}$. Following LSS, we use a flexible time-to-build technology:

$$i_t = \left(\sum_{s=0}^{S} \alpha_s^{\frac{1}{\sigma}} z_{t,t-s}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \tag{1}$$

where σ, $\alpha_s \geq 0$ and $\sum_s \alpha_s = 1$. In the extreme cases this would represent either a fixed-coefficient traditional time-to-build framework ($\sigma = 0$) or completely flexible, so that timing does not matter ($\sigma \to \infty$). For intermediate values the inputs will get spread throughout the $S+1$ time periods optimally as a smooth function of prices and the other fundamentals. Because only the final period input $z_{t,t}$ is chosen contemporaneously with i_t, the prior inputs represent state variables of the system.

There are several possible interpretations of (1). One is as mentioned, that each component $z_{t,t-s}$ represents a stage of production that must occur at date $t-s$ in order to produce i_t. The final good-producing firm could be thought of as either assembling capital itself, or simply purchasing capital that was assembled by another firm. In this case value-added is taking place at each stage. A second interpretation is that the combination of inputs for i_t all occurs at date t, but the various components must be ordered s periods in advance so that they will arrive at t. In either case the finite elasticity of substitution between the inputs represents some ability either to modify the timing or to slow down or speed up the deliveries.

In this paper we will adopt the second view. The supply chain literature and data emphasizes lead times for orders as the dominant factor in the dynamics of the production process. For example, the November 2015 report on manufacturing of the Institute for Supply Management shows an "Average commitment lead time for Capital Expenditures" of 135 days, with a fairly uniform distribution over periods ranging from "Hand-to-mouth" all the way to "1 year +." Of course the reality is a mixture: Some capital such as structures takes months or years to build, whereas some can be ordered to specification but with a waiting time for delivery. Also, behind the lead time is presumably some time-to-build for whatever is being ordered, in addition to queuing or backlogs.

At each date the firm chooses inputs for the investment at dates t through $t+S$, $\{z_{t,t}, z_{t+1,t}, ..., z_{t+S,t}\}$, which we will denote by $\tilde{z}_t$. Once chosen, all but $z_{t,t}$ become state variables as of date $t+1$, and will be denoted by $\{\bar{z}_{t+1,t}, ..., \bar{z}_{t+S,t}\}$. For example, if $S=1$, at date t the firm chooses $z_{t,t}$ and $z_{t+1,t}$ to determine i_t, k_{t+1}, $\bar{z}_{t+1,t}$, along with n_t given the wage w_t. Thus in general, under certainty the firm solves the problem

$$\max_{\tilde{z}_t, n_t, i_t, k_{t+1}} E_t \left\{ \sum_{t=0}^{\infty} \beta^t \left(\Lambda_t F(k_t, n_t) - \sum_{s=0}^{S} p_{t+s,t} z_{t+s,t} - \psi(i_t/k_t) k_t - w_t n_t \right) \right\}$$

subject to

$$q_t \; : \; k_{t+1} = k_t(1-\delta) + i_t \tag{2}$$

$$\mu_t \; : \; i_t = \left(\sum_{s=0}^{S} \alpha_s^{\frac{1}{\sigma}} z_{t,t-s}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \tag{3}$$

$$\lambda_{t+s,t} \; : \; \bar{z}_{t+s,t} = z_{t+s,t} \quad s = 1, ..., S. \tag{4}$$

We will use standard functional forms for F and ψ: $F(k,n) = k^{\phi} n^{1-\phi}$, and $\psi(i/k) = B(i/k - \delta)^{\eta}$, where B is some positive constant and $\eta > 1$. Λ_t represents Hicks-neutral technical progress, i.e. the technology shock process.

We can denote the S predetermined z variables as of time t as $\bar{z}_{t-1}$. The Lagrangian for this problem is

$$
\begin{aligned}
&\mathfrak{L}\left(k_t;\bar{z}_{t-1}\right) \\
= &\sum_{t=0}^{\infty}\beta^t E_t\left\{\Lambda_t F\left(k_t,n_t\right)-\sum_{s=0}^{S}p_{t+s,t}z_{t+s,t}-\psi\left(i_t/k_t\right)k_t-w_t n_t+q_t\left(k_t\left(1-\delta\right)+i_t-k_{t+1}\right)\right. \\
&\left.+\sum_{s=1}^{S}\lambda_{t+s,t}\left(\bar{z}_{t+s,t}-z_{t+s,t}\right)+\mu_t\left(\left(\alpha_0^{\frac{1}{\sigma}}z_{t,t}^{\frac{\sigma-1}{\sigma}}+\sum_{s=1}^{S}\alpha_s^{\frac{1}{\sigma}}\bar{z}_{t,t-s}^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}-i_t\right)\right\}.
\end{aligned}
$$

The first-order conditions for this problem are:

$$
\begin{aligned}
n_t &: w_t=\Lambda_t\left(1-\phi\right)\left(\frac{k_t}{n_t}\right)^{\phi} && (5)\\
k_{t+1} &: q_t=\beta E_t\left\{\Lambda_{t+1}\phi\left(\frac{k_{t+1}}{n_{t+1}}\right)^{\phi-1}-\right. && (6)\\
&\left.B\left(\left(\frac{i_t}{k_t}-\delta\right)^{\eta}-\eta\left(\frac{i_t}{k_t}-\delta\right)^{\eta-1}\frac{i_t}{k_t}\right)+\left(1-\delta\right)q_{t+1}\right\} && (7)\\
i_t &: q_t=\eta\left(\frac{i_t}{k_t}-\delta\right)^{\eta-1}+\mu_t && (8)\\
z_{t,t} &: p_{t,t}=\mu_t\left(\frac{z_{t,t}}{\alpha_0 i_t}\right)^{-1/\sigma} && (9)\\
z_{t+s,t} &: p_{t+s,t}=\lambda_{t+s,t}\quad s=1,...,S && (10)\\
\bar{z}_{t+s,t} &: \lambda_{t+s,t}=\beta^s E_t\left\{\mu_{t+s}\left(\frac{\bar{z}_{t+s,t}}{\alpha_s i_{t+s}}\right)^{-1/\sigma}\right\}\quad s=1,...,S. && (11)
\end{aligned}
$$

It is straightforward to eliminate the multipliers other than q_t, which replaces $(8)-(11)$ with

$$
\begin{aligned}
q_t &= \eta\left(\frac{i_t}{k_t}-\delta\right)^{\eta-1}+p_{t,t}\left(\frac{z_{t,t}}{\alpha_0 i_t}\right)^{1/\sigma} && (12)\\
1 &= \beta^s E_t\left\{\frac{p_{t+s,t+s}}{p_{t+s,t}}\left(\frac{\alpha_s z_{t+s,t+s}}{\alpha_0\bar{z}_{t+s,t}}\right)^{1/\sigma}\right\}\quad s=1,...,S. && (13)
\end{aligned}
$$

To complete the system, we also have an AR(2) process for Λ,

$$
\ln\left(\Lambda_t\right)=\rho_1\ln\left(\Lambda_{t-1}\right)+\rho_2\ln\left(\Lambda_{t-2}\right)+\epsilon_t \qquad (14)
$$

and the constraints $(2)-(4)$. (Note: Although this shock process can be "hump-shaped," by itself it does not produce a hump-shaped investment response in the standard q model.) In practice, because F exhibits constant returns to scale, we cannot pin down k and n separately, so we set $n=1$ and solve for w.

Under our interpretation, each of the elements of $\tilde{z}_t$ represent new orders. At the same time, at each date t, previous orders $\{\bar{z}_{t,t-S},...,\bar{z}_{t,t-1}\}$ along with $z_{t,t}$ are filled. If we let U_t denote the value of the stock of unfilled orders at the beginning of period t, we have

$$U_{t+1}=U_t+\sum_{s=1}^{S}p_{t+s,t}z_{t+s,t}-\sum_{s=1}^{S}p_{t,t-s}z_{t,t-s}. \tag{15}$$

which obeys the standard identity that the change in unfilled orders equals new orders minus shipments, or

$$\Delta U_{t+1}=no_t-y_t$$

where no_t denotes new orders at date t, and y_t denotes shipments. We can compute the level of unfilled orders directly from adding up all the orders from previous periods that have yet to be filled:

$$\begin{aligned} U_t &= \sum_{s=1}^{S}p_{t,t-s}z_{t,t-s}+\sum_{s=1}^{S-1}p_{t+1,t-s}z_{t+1,t-s}+\sum_{s=1}^{S-2}p_{t+2,t-s}z_{t+2,t-s}+... && (16)\\ &= \sum_{\tau=0}^{S-1}\sum_{s=1}^{S-\tau}p_{t+\tau,t-s}z_{t+\tau,t-s} && (17)\end{aligned}$$

We treat $z_{t,t}$ as a "hand-to-mouth" purchase, so it does not figure in unfilled orders. Unfilled orders thus includes all of the orders from previous periods

To see this more clearly in the simplest case with $S=1$, we have the firm choosing only

$z_{t,t}$ and $z_{t+1,t}$ given $\bar{z}_{t,t-1}$. The Lagrangian is

$$\begin{aligned}
& \mathfrak{L}\left(k_t; \bar{z}_{t,t-1}\right) \\
= & \sum_{t=0}^{\infty} \beta^t E_t \left\{ \Lambda_t F\left(k_t, n_t\right) - p_{t,t} z_{t,t} - p_{t+1,t} z_{t+1,t} - \psi\left(i_t/k_t\right) k_t - w_t n_t + \right. \\
& \left. q_t\left(k_t\left(1-\delta\right) + i_t - k_{t+1}\right) + \lambda_t\left(z_{t+1,t} - \bar{z}_{t+1,t}\right) + \mu_t\left(\left(\alpha_0^{\frac{1}{\sigma}} z_{t,t}^{\frac{\sigma-1}{\sigma}} + \alpha_1^{\frac{1}{\sigma}} \bar{z}_{t,t-1}^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} - i_t\right)\right\}
\end{aligned}$$

We assume that the economy is in steady state as of $t = 0$, with $\Lambda = 1$, $k_0 = k^*$, $\bar{z}_{t,t-1} = \alpha_1 \delta k^*$. Then $\epsilon_1 > 0$, and thereafter Λ follows (14)where $\rho \in [0, 1)$. So we have first-order conditions

$$\begin{aligned}
n_t &: w_t = \Lambda_t\left(1-\phi\right)\left(\frac{k_t}{n_t}\right)^{\phi} && (18) \\
k_{t+1} &: q_t = \beta E_t \left\{ \Lambda_{t+1}\phi\left(\frac{k_{t+1}}{n_{t+1}}\right)^{\phi-1} \right. && (19) \\
& \left. -B\left(\left(\frac{i_t}{k_t} - \delta\right)^{\eta} - \eta\left(\frac{i_t}{k_t} - \delta\right)^{\eta-1}\frac{i_t}{k_t}\right) + \left(1-\delta\right) q_{t+1} \right\} && (20) \\
i_t &: q_t = \eta\left(\frac{i_t}{k_t} - \delta\right)^{\eta-1} + p_{t,t}\left(\frac{z_{t,t}}{\alpha_0 i_t}\right)^{1/\sigma} && (21) \\
z_{t+1,t} &: 1 = \beta E_t \left\{ \frac{p_{t+1,t+1}}{p_{t+1,t}}\left(\frac{\alpha_s z_{t+1,t+1}}{\alpha_0 \bar{z}_{t+1,t}}\right)^{1/\sigma} \right\} && (22)
\end{aligned}$$

and from the constraints, we know that

$$\begin{aligned}
z_{t,t} &= \left(\frac{i_t^{\frac{\sigma-1}{\sigma}} - \alpha_1^{\frac{1}{\sigma}} \bar{z}_{t,t-1}^{\frac{\sigma-1}{\sigma}}}{\alpha_0^{\frac{1}{\sigma}}}\right)^{\frac{\sigma}{\sigma-1}} \\
k_{t+1} &= k_t\left(1-\delta\right) + i_t.
\end{aligned}$$

Finally, we have

$$\begin{aligned}
U_t &= p_{t,t-1} z_{t,t-1} \\
U_{t+1} &= U_t + p_{t+1,t} z_{t+1,t} - p_{t,t-1} z_{t,t-1}.
\end{aligned}$$

This of course is an almost trivial case, as each unfilled order is shipped the next period, so the ratio of unfilled orders to shipments is always one. When $S > 1$, U_t will contain a mixture of orders made in different time periods timed to arrive at different dates, and some components of U_t will remain in U_{t+1}.

For quarterly data, the ISM report alluded to above (as well as shipments and orders data discussed below) suggests that $S = 5$ is reasonable. Figure 1 displays impulse responses of i and q from a persistent shock to Λ, with $\rho_1 = 1.4$, $\rho_2 = -0.45$, $S = 5$, and the standard deviation of the shock set to match the volatility of shipments (as described in the next section). For simplicity we do not model production of the z inputs, and just set their prices equal to one. We see that q responds similarly to the standard model without delivery lags, but—not surprisingly—investment starts low and builds up over four periods before turning back down toward its long-run steady state value.

Figure 2 displays the behavior of unfilled orders, and of the ratio of unfilled orders to shipments (total deliveries of the z inputs at date t except for $z_{t,t}$, which is ordered and delivered within the period), denoted by y. Unfilled orders behaves similar to investment, though smoothed, but relative to shipments it is hump-shaped with a return to very close (in fact slighly below) the steady state after four periods.

Figure 3 compares the response of new orders and shipments. Not surprisingly, given the planned increases in investment depicted in Figure 1, new orders jump immediately. Shipments of course mainly reflect past decisions, and therefore have a more muted immediate response, followed by a buildup that is similar to the pattern in investment (as it must be).

2 Data on Orders and Shipments

The model in the previous section was developed with a view to generating the hump-shaped response of investment to shocks that has been observed in data, and that contradicts simple adjustment cost models such as the now-classic q model (e.g. Hayashi, 1982). That being the case, we will not treat Figure 1 as any kind of confirmation of the approach in this paper, only as a prerequisite to further consideration of the model. Support for the model must come from its additional implications. Consequently in this section we examine the data on new orders, shipments, and unfilled orders.

Before proceeding to that we will mention one other aspect of the investment-q relationship: Related to the hump-shaped response is the fact that in regressions of investment (or i/k) on q the coefficient is typically implausibly small (indicative of implausibly large adjustment costs). While the model thus far is too simple to quantify, it is obvious qualitatively that the relationship between investment and q is weaker the higher is S. As a rough guide, the numerical results from the model suggest that the approximate regression coefficient of i on q declines from about 1.94 to 0.82 as S goes from 1 to 5.

Now to the data: The Commerce Department releases monthly data on new orders, unfilled orders, and shipments from the manufacturing sector. To match the model's frequency we aggregate those data up to the quarterly frequency. As the series are available only in dollars, we deflate by the PPI for capital goods.[1] Figure 4 depicts new orders, shipments, and the stock of unfilled orders for capital goods, seasonally adjusted, quarterly from 1992:Q2-2015Q3. We see that all three are markedly pro-cyclical, with the amplitude of new orders noticeably larger than that of shipments. The standard deviation of detrended log new orders is 0.126, versus 0.085 for shipments. Figure 5 displays the ratio of unfilled orders (as of the end of the previous time period) to shipments, with and without detrending. The ratio averages around 2.4 quarters (consistent, as alluded to earlier, with the assumption of $S = 5$), though clearly drifts over time. It also clearly lags the cycle by several quarters, to the point that it almost looks counter-cyclical. What seems to happen is that new orders begin to drop precipitously at the onset of a recession, whereas unfilled orders decline more slowly because of the momentum built in from previous periods' orders, so the ratio actually moves counter to the cycle at first.

[1] In general this is problematic for a variety of reasons (see Nalwaik and Pinto, 2015, for example). For our purposes at the moment the manner of deflation makes little difference.

We see these characteristics in the behavior of the model, albeit with somewhat less inertia than in the data. Note that in Figure 2, in response to a positive technology shock, the ratio actually declines at first (falling below the steady state level) before rising. And in Figure 3 we see new orders jumping ahead of shipments, with the latter rising gradually, only catching up after about a year. The model also implies that the standard deviation of shipments exceeds that of orders, though by quantitatively much less than in the data. Specifically, we get a standard deviation for shipments (relative to the mean) of 0.083, and for new orders of 0.09.

3 Conclusions and Plans for Further Work

We have developed a flexible "time-to-plan" model, based on Lubik et al. (2015), for investment that focuses on the supply side of the capital goods market. Relative to the practice of assuming ad hoc investment adjustment costs it has legitimate microeconomic foundations, and is able to replicate qualitative features of not only investment data, but also data on new orders, shipments, and unfilled orders.

Thus far we have not modeled the production of capital goods, and simply assumed an infinitely elastic supply curve for the z inputs at a unit price. If these goods are produced with short-run increasing marginal cost, then their prices will increase with demand and thereby induce some further spreading of investment toward the future. Specifically, in response to an aggregate shock such as that to Λ_t, the prices of the inputs will rise inversely to the lead time, i.e. $z_{t,t}$ will increase the most, followed by $z_{t+1,t}$ and so on, causing some corresponding substitution. We also plan to look at industry–level data, where heterogeneity in lead times should have predictions for the dynamics of investments, shipments, and orders. This heterogeneity may also be important in understanding the behavior of aggregates.

Another limitation of the model is that the maximum lead time is fixed by the parameter S. This tends to limit the cyclicality of unfilled orders and hence the relative volatility of new orders compared to shipments. We hope to explore ways to allow this to vary endogenously, though to do so in this framework will be challenging.

Finally, it will also be important to extend the model to general equilibrium, as in Lubik et al. This is clearly feasible. It would involve production in both consumption and capital goods sectors, would endogenize interest rates and other prices, and allow consideration of other types of shocks, and yield quantitative implications that can be compared with business cycle fluctuations.

References

[1] Basu, S. and M. Kimball, 2003, "Investment Planning and the Effects of Fiscal and Monetary Policy" manuscript.

[2] Carlton, D., 1983. "Equilibrium Fluctuations when Price and Delivery Lag Clear the Market," *The Bell Journal of Economics* 14(2), 562-572.

[3] Chari, V. V., P. Kehoe, and E. McGratton, 2000, "Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem," *Econometrica*, 68 (5), 1151-1179.

[4] Casares, M., 2006, "Time-to-Build, Monetary Shocks, and Aggregate Fluctuations," *Journal of Monetary Economics,* 53*,* 1161-1176.

[5] Del Boca, A., M. Galeotti, C. Himmelberg, and P. Rota, 2010, "Investment and Time to Plan and Build: A Comparison of Structures vs. Equipment in a Panel of Italian Firms," *Journal of the European Economic Association,* 6(4), 864-869.

[6] Christiano, L., M. Eichenbaum, and C. Evans, 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy* 113(1), 1-45.

[7] Edge, R. M., 2007, "Time-to-Build, Time-to-Plan, Habit-Persistence, and the Liquidity Effect," *Journal of Monetary Economics,* 54*,* 1644-1669.

[8] Haltiwanger, J. and L. Maccini, 1989, "Inventories, Orders, Temporary and Permaqnent Layoffs: An Econometric Analysis," *Carnegie-Rochester Conference Series on Public Policy,* 30, 301-366.

[9] Hayashi, F., 1982. "Tobin's Marginal q and Average q: A Neoclassical Interpretation," *Econometrica* 50(1), 213-224.

[10] Institute for Supply Management, December 2015. *November Manufacturing ISM Report on Business*, https://www.instituteforsupplymanagement.org/ISMReport/MfgROB.cfm

[11] Kahn, J., 2010, "The Great Moderation and the Great Recession," manuscript.

[12] Kydland, F. E. and E. C. Prescott, 1982, "Time to Build and Aggregate Fluctuations", *Econometrica,* 55, 1345-1370.

[13] Lubik, T., J. Sarte, and F. Schwartzman, 2015. "What Inventory Behavior Tells Us About How Business Cycles Have Changed", manuscript.

[14] Lucca, D. O., 2007, "Resuscitating Time-to-Build," manuscript.

[15] Maccini, L. 1973. "On Optimal Delivery Lags," *Journal of Economic Theory* 6, 107-125.

[16] Nalewaik, J. J., and E. P. Pinto, 2015, "The Response of Capital Goods Shipments over the Business Cycle," *Journal of Macroeconomics* 43, 62-80.

[17] Reagan, P. and D.P. Sheehan, 1985, "The Stylized Facts about the Behavior of Manufacturers' Inventories and Backorders over the Business Cycle: 1959-1980," *Journal of Monetary Economics,* 15*,* 214-246.

[18] Tsoukalas, J., 2011, "Time to Build Capital: Revisiting Investment Cash-flow Sensitivities", *Journal of Economic Dynamics and Control,* 35, 100-1016.

[19] West, K., 1988, "Order Backlogs and Production Smoothing", in *The Economics of Inventory Management,* A. Chikan and M. Lovell (eds.), (Amsterdam: Elsevier), 305-317; reprinted in T. Kollintzas (ed.), *The Rational Expectations Inventory Model,* (New York: Springer-Verlag), 1989, 246-279..

Table 1: Parameters and Steady StateParameters ($S = 5$)

Parameters							
α_s	β	σ	ρ_1	ρ_2	η	δ	ϕ
1/6	0.99	2	1.4	0.45	2	0.1	0.33

Steady State Values							
k	i	z_s	q	U	U/y	w	no, y
5.15	0.515	0.086	1	1.24	2.47	1.15	0.502

Figure 1: Impulse Responses of Investment and q

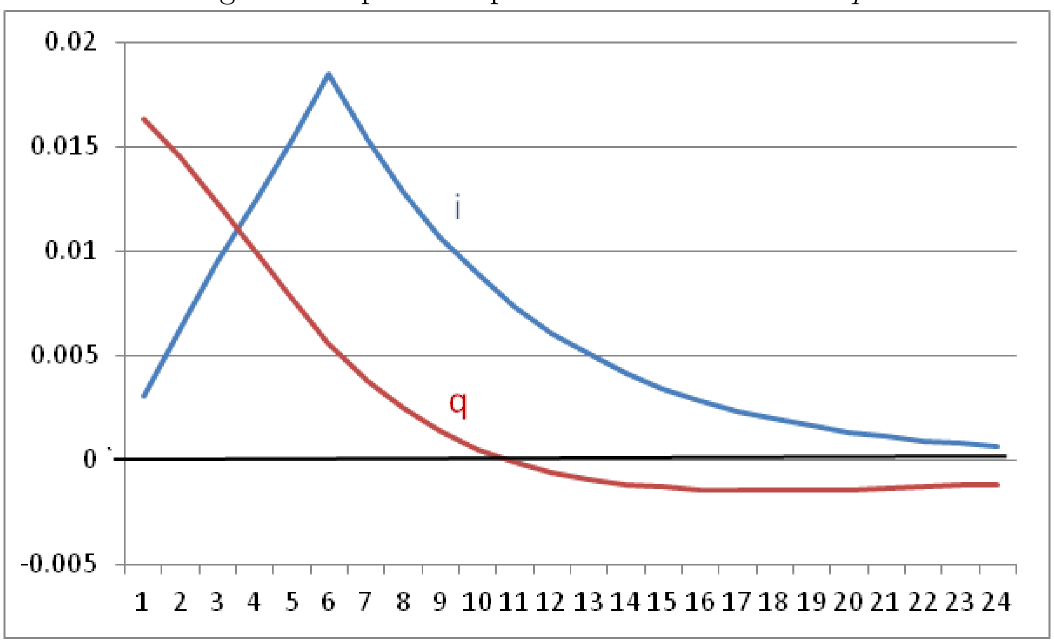

Figure 2: Response of Unfilled Orders and Ratio of Unfilled Orders to Shipments

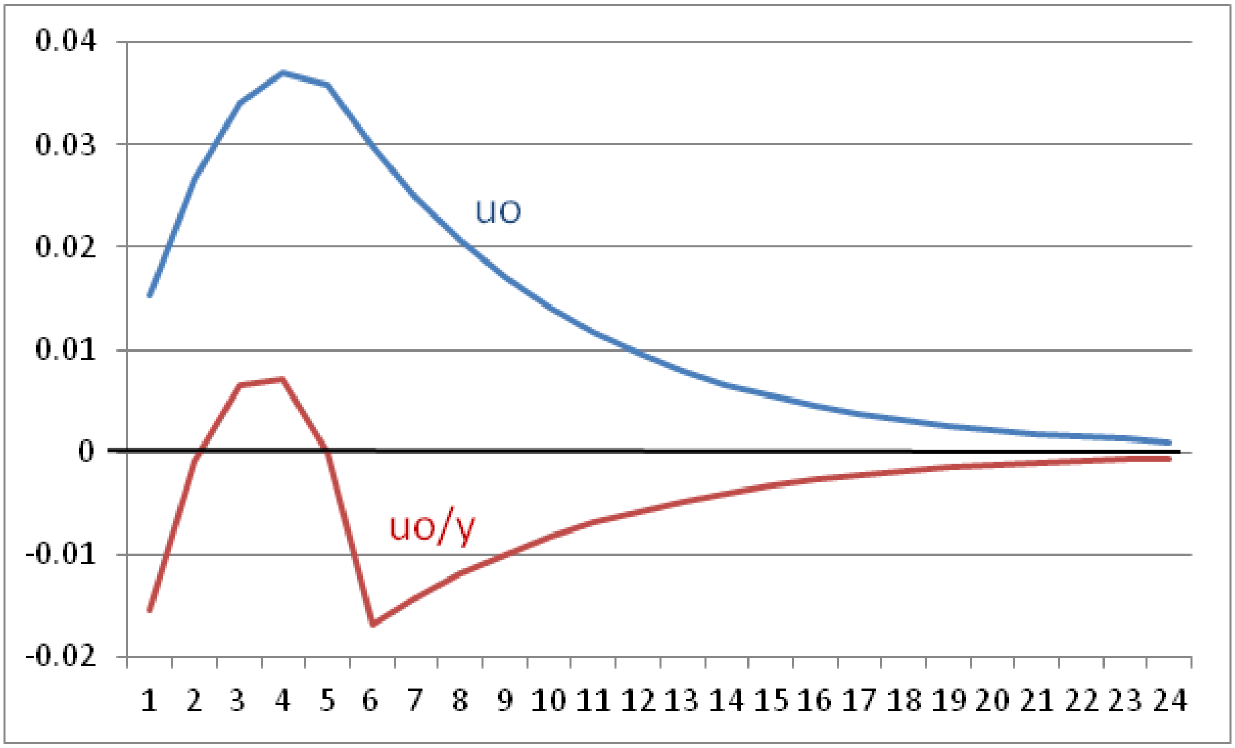

Figure 3: Response of New Orders and Shipments

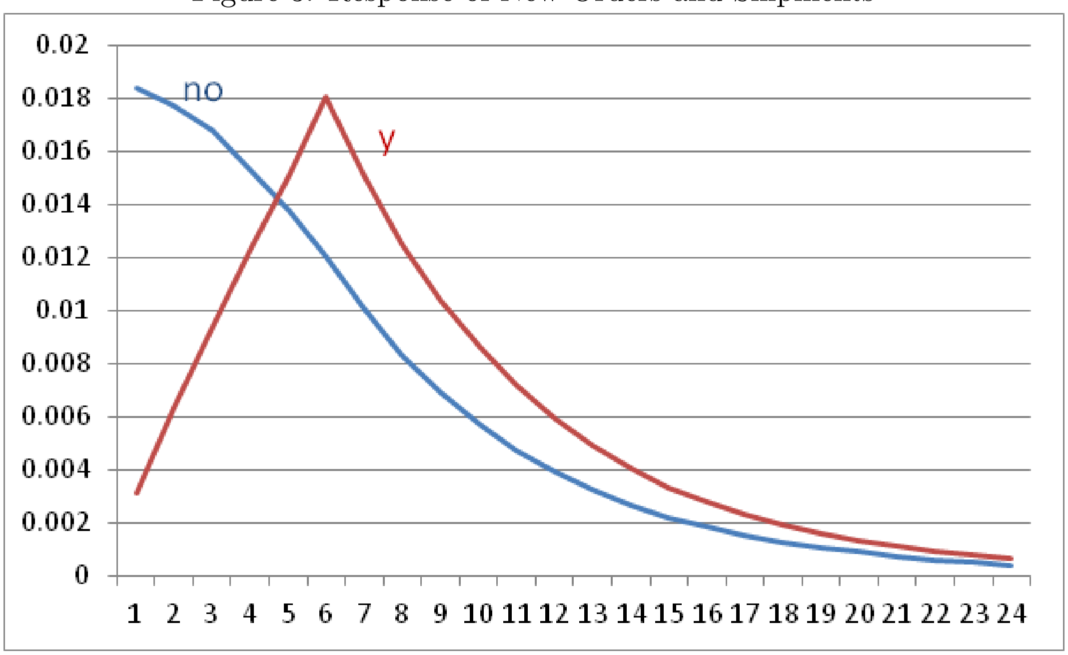

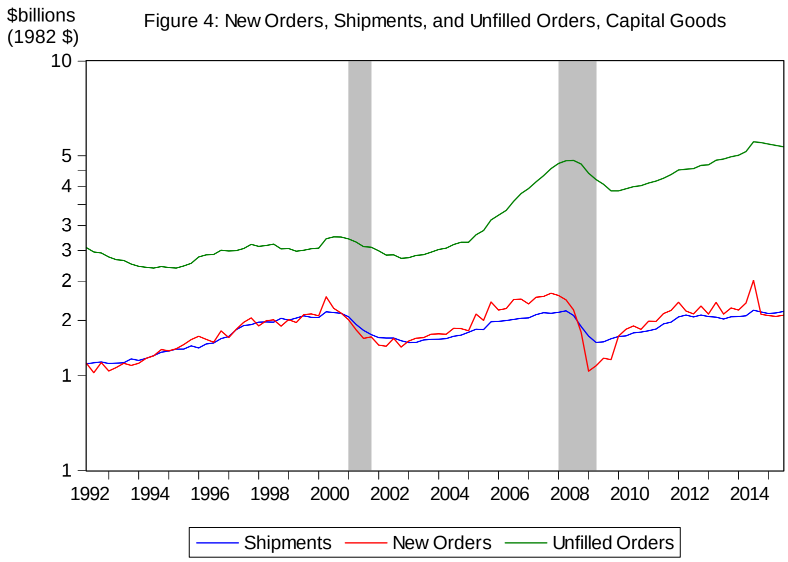

Shaded periods denote NBER-defined recessions

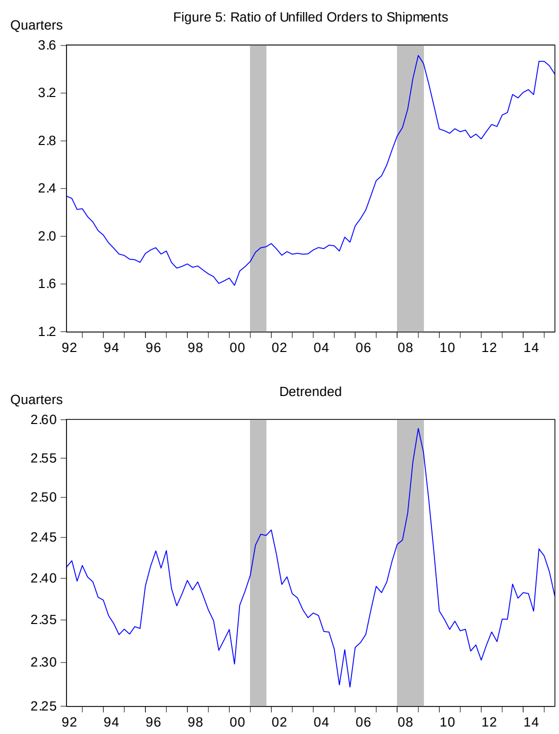

Shaded periods denote NBER-defined recessions