TURKISH ECONOMIC ASSOCIATION

DISCUSSION PAPER 2004/15
http://www.tek.org.tr

PARTIAL ADJUSTMENT WITHOUT APOLOGY

Robert G. King and Julia K. Thomas

June, 2004

Partial adjustment without apology

Robert G. King

Boston University, Federal Reserve Bank of Richmond and NBER

Julia K. Thomas[1]

University of Minnesota, Federal Reserve Bank of Minneapolis and NBER

March 01, 2004

[1]Please direct correspondence to Julia Thomas, Department of Economics, University of Minnesota, 271 19th Avenue South, Minneapolis, MN 55455; tel: 612 626 9675; email: jkt@econ.umn.edu. We thank seminar participants at the Oxford Meetings of the Society of Economic Dynamics and the University of Virginia for their comments. All remaining errors are our sole responsibility. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Richmond or Minneapolis, or the Federal Reserve System.

Abstract

Many kinds of economic behavior appear to be governed by discrete and occasional individual choices. Despite this, econometric partial adjustment models perform relatively well at the aggregate level. Analyzing the classic employment adjustment problem, we show how discrete and occasional microeconomic adjustment is well described by a new form of partial adjustment model that aggregates the actions of a large number of heterogeneous producers.

We begin by describing a basic model of discrete and occasional adjustment at the micro level, where production units are essentially restricted to either operate with a fixed number of workers or shut down. We show that this simple model is observationally equivalent at the market level to the standard rational expectations partial adjustment model. We then construct a related, but more realistic, model that incorporates the idea that increases or decreases in the size of an establishment's workforce are subject to fixed adjustment costs. In the market equilibrium of this model, employment responses to aggregate disturbances include changes both in employment selected by individual establishments and in the measure of establishments actively undertaking adjustment. Yet the model retains a partial adjustment flavor in its aggregate responses. Moreover, in contrast to existing models of discrete adjustment, our *generalized partial adjustment model* is sufficiently tractable to allow extension to general equilibrium.

1 Introduction

In many contexts, actual factor demands clearly involve complicated dynamic elements absent in static demand theory. For example, empirical studies of the market demand for labor typically find that lags, either of demand or of the determinants of demand, contribute substantially to the explanation of employment determination. The most frequent rationalization of such lags is that individual plants face marginal costs of adjustment that are increasing in extent of adjustment, leading them to choose *partial adjustment* toward the levels suggested by static demand theory.

Many empirical studies also indicate, however, that the partial adjustment model is inconsistent with the behavior of individual plants or firms. For example, Hamermesh (1989) shows that individual plants undertake discrete and occasional workforce adjustments rather than the smooth changes implied by partial adjustment. Nonetheless, the model continues to be a vehicle for applied work, essentially because it is a tractable way of capturing some important dynamic aspects of market demand. It is frequently thus employed in an apologetic manner, with the researcher suggesting that it is a description of market, rather than individual, factor demand.[1] In fact, Hamermesh (1989) finds that a labor demand aggregate, made up of just seven plants, appears at least as well described by the partial adjustment model as by positing a representative firm that adjusts in a discrete and occasional manner.

In this paper, we develop several models which embody the idea that individual production units adjust in a discrete and occasional manner, yet have the property that there is smooth adjustment at the aggregate level. In the first of these models, there is an exact observational equivalence at the aggregate level between the standard rational expectations

[1] See, for example, Kollintzas (1985).

model of partial adjustment (Sargent (1978)) and our model of an industry's employment demand. More generally, this model illustrates key features of our modeling approach. Specifically, individual units face differing fixed costs of adjustment, so that the timing of their adjustments is occasional and asynchronized. Nevertheless, aggregation across plants leads to a smooth pattern of industry labor demand that is well-approximated by the standard partial adjustment model. Thus, the structural features that result in gradual adjustment (distributed lags) also imply that individual plants base their employment decisions on expectations of future wages and productivities (distributed leads).

Our subsequent models provide a microeconomic foundation for the variety of plant-level adjustment examined in the empirical work of Caballero and Engel (1992, 1993) and Caballero, Engel, and Haltiwanger (1997). There, individual production units are assumed to adjust employment probabilistically, and that the probability of adjustment is a function of the difference between a target level of employment and actual employment. Aggregating from such adjustment hazard functions, which are their basic unit of analysis, they examine the implications of the resulting state-dependent adjustment behavior for aggregate employment demand dynamics. In the absence of a microeconomic foundation for such probabilistic adjustment, Caballero and Engel (1993, p. 360, paragraph 2) explain that they "trade some deep parameters for empirical richness." In contrast, we explicitly model the plant's adjustment decision as a *generalized* (S, s) problem and derive the adjustment hazard functions that are the starting point of previous research.[2]

We summarize some of the key stylized facts uncovered in the empirical literature and require that our theoretical models be consistent with them. One such finding is that an important route by which aggregate shocks affect aggregate employment is by changing the

[2] Generalized (S, s) models were first studied by Caballero and Engel (1999) to explain the observed lumpiness of plant-level investment demand.

fraction of plants that choose to adjust. Accordingly, we move from our initial model to develop a generalized partial adjustment model in which the aggregate adjustment rate is an endogneous function of the state of the economy. In doing so, we relax the observational equivalence to the traditional partial adjustment model where aggregate adjustment rates are invariant to changes in economic policy. Nonetheless, impulse responses establish that our generalized model retains the basic features of gradual partial adjustment. Another distinguishing feature of our theoretical approach is that it is feasible to undertake generalized (S, s) analysis within a general equilibrium framework, so that the influence of aggregate shocks on equilibrium adjustment patterns may be systematically studied.

The organization of this discussion is as follows. In section 2, we present the standard partial adjustment model, and in section 3 we discuss the evidence on microeconomic adjustment patterns that this standard model fails to reproduce. In section 4, we present the simplest form of a gradual adjustment market labor demand model consistent with discrete and occasional adjustment choice at the plant level. We show that this model is observationally equivalent to the standard partial adjustment model, given suitable choice of parameters. However, this preliminary model is limited in that individual production units undertake only a single adjustment decision, which corresponds to entering productive activity with a fixed level of employment.

In section 5, building on the model of entry described above, we develop a model that is consistent with the observation that plants hire varying amounts of labor, with these quantities adjusted at discrete and occasional times. We use this model to illustrate a *hedging effect* on the demand for labor that arises when a plant recognizes that there may be future departures from the workforce prior to its next employment adjustment. Next, in section 5.1, we endogenize the timing of employment changes by assuming that each plant faces a fixed cost of adjustment that is random across both time and plants. The resulting generalized (S, s) model allows us to examine the influence of deep parameters on the adjustment process. Moreover, with a large number of plants, the model is similar to the traditional partial adjustment model in that it yields a smooth market labor demand. We illustrate the properties of this *generalized partial adjustment* model using a series of numerical examples.[3] A distinguishing feature of our model, over and above its consistency with the microeconomic evidence on employment adjustment, is that it is able to reproduce the sharp changes in market employment demand found in the data during episodes involving large changes in productivity.[4] We also illustrate how our framework may be tractably embedded within a fully specified general equilibrium macroeconomic model. Section 6 provides a brief conclusion.

2 The standard partial adjustment model

The standard partial adjustment model relates current employment, N_t, to *target* or desired employment, N_t^*, through

$$N_t - N_{t-1} = \kappa[N_t^* - N_{t-1}], \tag{1}$$

[3] Our generalized partial adjustment model is distinguished from earlier generalized cost of adjustment models, as summarized, extended and critiqued in Mortensen (1973), in that it suggests very different dynamics at the establishment-level. Nonetheless, because the final model that we present is essentially one with many dynamically related factor demands, it is capable of generating some of the aggregate dynamics that motivated researchers in this earlier area. For example, under unrestricted parameters, interrelated factor demand models were found to be consistent with oscillatory approaches to the long-run position. The generalized stochastic adjustment model that we develop can also generate such rich dynamics, although it does not do so under the parameters selected here.

[4] This is because the economywide rate of adjustment implied by our model varies with aggregate conditions. The traditional model under-predicts employment changes during such episodes precisely because the adjustment rate there is constant.

with $0 < \kappa < 1$ being the fraction of the gap $N_t^* - N_{t-1}$ that is closed in the period. This specification implies the influence of past actual or desired employment on current employment,

$$N_t = \kappa N_t^* + (1-\kappa)N_{t-1} = \kappa \sum_{j=0}^{\infty}(1-\kappa)^j N_{t-j}^*. \tag{2}$$

As shown by Sargent (1978), this empirical partial adjustment model may be derived as the solution to a firm's dynamic profit maximization problem under the assumption that there are quadratic costs of adjusting the workforce. To develop this standard partial adjustment model, assume that the firm's workforce declines, due to quits or mismatches, at the rate $d \in [0,1)$, in the absence of any costly employment-adjusting action. If e_t is the number of employees hired at time t, $t = 0, \ldots$, then

$$N_t = (1-d)\, N_{t-1} + e_t \tag{3}$$

and the cost of actively adjusting the workforce is $\Xi\,(e_t)$. The quadratic cost assumption is that

$$\Xi\,(e_t) = \frac{B}{2} e_t^2, \tag{4}$$

where $B > 0$ is a cost parameter. Equation (4) captures the idea that the firm's marginal adjustment cost is rising in the extent of employment adjustment.[5]

Let A_t represent a productivity shift term, and let W_t represent the real wage at time t, $t = 0, 1, \ldots$. Both A_t and W_t are serially correlated random variables, known at the beginning of period t. The flow profit of the firm at time t, π_t, is output $f(N_t, A_t)$ less adjustment costs, $\Xi\,(e_t)$, and the wage bill, $W_t N_t$.

$$\pi_t = f(N_t, A_t) - \Xi\,(e_t) - W_t N_t.$$

[5] This same idea is incorporated in alternative adjustment cost functions that are used in applied work.

Discounting future earnings by the constant factor, $\beta \in (0,1)$, the firm solves the following optimization problem:

$$\max_{\{N_t,e_t\}_{t=0}^{\infty}} \mathbf{E}\left(\sum_{t=0}^{\infty} \beta^t \Big[f(N_t, A_t) - \Xi\left(e_t\right) - W_t N_t \Big] \Bigg| \left(A_0, W_0\right) \right)$$

$$\text{subject to (3), } N_{-1} \text{ given.}$$

Let v_t represent the current-value multiplier associated with (3). Then the efficient choice of labor requires:

$$\frac{\partial f(N_t, A_t)}{\partial N_t} = W_t + v_t - \beta \mathbf{E}_t v_{t+1} \left(1-d\right), \tag{5}$$

while the corresponding condition for gross employment changes, e_t, is

$$v_t = \frac{\partial \Xi\left(e_t\right)}{\partial e_t}. \tag{6}$$

Using (3), (4) and (6) to simplify (5), we have

$$\frac{\partial f(N_t, A_t)}{\partial N_t} = W_t + B\left(N_t - \left(1-d\right) N_{t-1}\right) - \beta\left(1-d\right) \mathbf{E}_t B\left(N_{t+1} - \left(1-d\right) N_t\right). \tag{7}$$

Assuming that the production function is quadratic or, more generally, approximating it using a second-order Taylor expansion,

$$f(N_t, A_t) \equiv f + f_n N_t + \frac{1}{2} f_{nn} N_t^2 + f_{na} A_t N_t + f_a A_t + \frac{1}{2} f_{aa} A_t^2, \tag{8}$$

and defining $\Phi = \frac{B - f_{nn} + \beta(1-d)^2 B}{(1-d)B}$, we can rewrite (7) as

$$\beta \mathbf{E}_t\left(N_{t+1}\right) - \Phi N_t + N_{t-1} = \frac{W_t - f_{na} A_t - f_n}{\left(1-d\right) B}. \tag{9}$$

Note that $\Phi > 0$, since $f_{nn} < 0$ is required by concavity of the production function. This is sufficient to ensure that the second-order stochastic difference equation (9) has two real roots, $\mu_1 \in (0, 1-d)$ and $\mu_2 > [\beta(1-d)]^{-1}$, that jointly solve

$$\begin{aligned}\mu_1 \mu_2 &= \frac{1}{\beta} \\ \mu_1 + \mu_2 &= \frac{\Phi}{\beta}.\end{aligned}$$

The firm's target employment in (1) is then given by

$$N_t^* = \frac{\mu_1}{B\,(1-d)\,(1-\mu_1)} \left(\mathbf{E}\Big[\sum_{j=0}^{\infty} (\frac{1}{\mu_2})^j \left(f_{na} A_{t+j} - W_{t+j} \right) \Big| (A_t, W_t) \Big] + \frac{f_n}{1 - \frac{1}{\mu_2}} \right), \tag{10}$$

and its adjustment rate is

$$\kappa = 1 - \mu_1. \tag{11}$$

Our expression for target employment illustrates Sargent's (1978) result that the presence of lags in employment, as in (1) under rational expectations, implies leads. Expectations of future wages and productivity influence the current employment target since its choice, given adjustment costs, will in part determine future employment. This presence of expectational leads dampens the response of current employment to changes in current wage and productivity and yields smooth, gradual changes in employment over time. Current employment, N_t, is directly related to lagged, N_{t-1}. Moreover, the other determinant of current employment, target employment N_t^*, is a discounted sum of future wages and productivities, and, as such, is only partly determined by current wage and productivity.

3 Disconcerting evidence

While the traditional partial adjustment model offers a tractable framework within which to study gradual aggregate labor adjustment, there is considerable empirical evidence to suggest that the model is *not* consistent with the behavior of individual production units. This evidence also suggests a number of stylized facts about individual and aggregate adjustment, which this section summarizes.

Stylized fact 1: Adjustment at the plant level is discrete, occasional and asynchronous. Hamermesh (1989) examines monthly data on output and employment between 1983 and 1987 across seven manufacturing plants. For each plant, output fluctuates substantially over the sample. Employment exhibits long periods of constancy broken by infrequent, but large, jumps at times roughly coinciding with the largest output fluctuations. Hence, the plant data are not consistent with the smooth employment adjustment that would arise from convex adjustment costs.

Stylized fact 2: Aggregates exhibit smooth and partial adjustment. Hamermesh (1989) also examines the behavior of the aggregate of his seven manufacturing plants. He finds that fluctuations in aggregate employment across plants resembles the dynamics of aggregate output and appears consistent with smooth adjustment behavior of aggregates. More specifically, Hamermesh argues that the standard partial adjustment model works quite well at the aggregate level, even though it does not describe the behavior of individual production units.[6]

[6] In particular, he compares log likelihood values from the estimation of a smooth adjustment model based on quadratic adjustment costs with those from a lumpy adjustment fixed-cost alternative. For plant level data, the latter model achieves much larger likelihood values, indicating that lumpy adjustment based on fixed costs better describes the plant level data. Further, the switching model estimates of the percentage 'disequilibrium' required to induce adjustment are large. This indicates that plants vary employment with

Stylized fact 3: Adjustment hazards depend on aggregate conditions. Following the econometric literature on discrete choices, the probability that an individual production unit makes a discrete change during a particular time period is typically called an *adjustment hazard* in the literature. Caballero and Engel (1993) construct a general framework for studying aggregate employment changes that can incorporate a variety of assumptions about how adjustment hazards at the level of the individual production unit are related to aggregate conditions. Using data on U.S. manufacturing employment from 1961 through 1983, Caballero and Engel examine the dynamics of aggregate employment changes under two alternative specifications for the hazard function: (1) a benchmark constant hazard case and (2) an alternative hazard model involving higher moments of the cross-sectional distribution of firms' 'disequilibrium' levels, representing state-dependent adjustment behavior. They find large increases in explanatory power for aggregate employment changes in moving from the constant hazard model to a generalized hazard structure and attribute this to the effects of large aggregate shocks upon the employment hazard.

Stylized fact 4: Adjustment hazards depend on measures of 'micro gaps'. More direct evidence on the importance of state-dependent adjustment hazards is provided by Caballero, Engel and Haltiwanger (1997). Studying the direct relationship between the adjustment hazard at the level of the individual production unit and the extent of that unit's gap between current employment and a measure of desired employment, these authors show that

---

a non-marginal adjustment only in the presence of substantial shocks to expected output. However, the difference in the aggregate study is too small to discriminate between the two models, as is the case when the two models are compared using 4-digit SIC data. Thus, lumpy adjustment behavior at the microeconomic level is obscured by aggregation. This, and similar evidence, leads Hamermesh and Pfann (1996, page 1274) to conclude that "observing smooth adjustment based on data describing industries or higher aggregates over time is uninformative about firms' structures of adjustment costs and in no way disproves the existence of lumpy costs."

the adjustment hazard depends on the size of this discrepancy. They suggest that individual units may face differential adjustment costs, so that the distribution of adjustment costs governs the adjustment hazard.

Stylized fact 5: Aggregate shocks are much more important in accounting for aggregate responses than are shifts in cost distributions. The empirical analysis of Caballero, Engel and Haltiwanger (1997) also suggests that changes in the distribution of adjustment costs are not central in explaining stylized fact 3. Rather, aggregate shocks induce changes in hazards that are important for aggregates because they produce movements along the micro-distribution of employment imbalances.

These five facts motivate the development of the models in the balance of this paper.

4 Reinterpreting the partial adjustment model

We now develop a simple model that is designed to capture the three key findings in Hamermesh (1989), which we summarize as follows. First, at the level of individual production units, there is discrete and occasional adjustment. Second, heterogeneity in the circumstances of individual units leads these actors to adjust at different times. Third, the industry adjustment process is well-approximated by the standard partial adjustment model.

More specifically, our model of an industry's labor demand relies on a continuum of small production units in the industry, each of which faces a different fixed cost of activating production – an action which we call "entry" – by hiring a single unit of labor. We incorporate various assumptions to ensure tractable aggregation. Yet, while simple, our model economy will be exactly observationally equivalent to the standard partial adjustment model at the industry level.[7] Further, the central device – differential fixed costs

[7] To emphasize this equivalence, we use notation that is consistent with that introduced for the partial

of adjustment – can be used in richer models that have many features absent in our example, including various margins of adjustment at the level of individual actors and labor market equilibrium or general equilibrium. These richer models can display classic partial adjustment behavior when individual actors face a distribution of fixed costs.

To begin, we assume that in each period there is a continuum of production locations, indexed by $z \in \mathbf{R}_+$. Each of these units may produce a final good using a technology that is identical across units. This production function is constant returns to scale in two factors, labor and an industry-specific factor which might be land. Labor at the z^{th} production unit is $n_t(z)$, while the amount of the other factor is $x_t(z)$. We assume that this industry-specific factor is in exogenous supply of X_t at the industry level. We also assume that it may be costlessly reallocated across locations. As before, the production function is affected by a productivity shifter A_t, which now is common to all locations. The production function for the z-th production unit is therefore written as

$$F(n_t(z), x_t(z), A_t). \tag{12}$$

Given constant returns-to-scale production, there is no natural limit on the scale of an active production unit. Coupling this with fixed costs of activating each unit, efficiency would dictate that all production be concentrated at a single location. Accordingly, following Parente and Prescott (1994), we assume that there is an upper bound on the amount of labor input that can be used at any single production location. For notational convenience, we assume that this limit is one unit. Calling the industry level labor demand N_t, industry production is $N_t F(1, x_t, A_t)$ with $x_t = X_t/N_t$. Equivalently, the industry has a production function

$$F(N_t, X_t, A_t) \tag{13}$$

adjustment model of section 2.

that exhibits decreasing returns due to the presence of the industry-specific factor.

Each period begins with a unit measure of potential production locations, each of which may be activated by a costly process that requires the payment of a fixed cost. Furthermore, there are N_{t-1} active production units that have previously paid fixed costs so as to enter into the industry. However, we assume that a fraction $d \geq 0$ of these are now randomly required to pay a new fixed cost; otherwise, they must exit. Thus, the stock of active production units evolves according to

$$N_t - N_{t-1} = e_t - dN_{t-1}$$

where e_t is the fraction of potential production locations that are activated in the current period through the payment of the fixed entry cost.

Fixed entry costs differ across production units, but are distributed on the unit interval, with a cumulative distribution function $G(\xi)$ where ξ is the fixed cost of an individual production unit. There is an associated probability density $g(\xi)$, so that, if fraction e of production units pays the fixed cost in a particular period, then the total adjustment cost of the industry is

$$\Xi(e) = \int_0^{G^{-1}(e)} \xi g(\xi) d\xi. \tag{14}$$

As an example, suppose that the CDF is $\mathbf{G}(\boldsymbol{\xi}) \equiv \xi/B$, so that costs are uniformly distributed from zero to a maximum value of B. Then, $g(\xi) = 1/B$ and

$$\Xi(e) = \frac{B}{2}e^2, \tag{15}$$

so that there is an industry adjustment cost which is quadratic in entry.

4.1 The industry equilibrium

We analyze the industry equilibrium, as in Lucas and Prescott (1971) and many subsequent papers, by first studying the optimal adjustment pattern in this section and then showing that this is the same as a competitive equilibrium allocation. The optimal allocation maximizes the expected present discounted value of industry profits, less costs of adjustment:

$$\max_{\{N_t,,X_t,e_t\}_{t=0}^{\infty}} \mathbf{E}\left(\sum_{t=0}^{\infty}\beta^t\Big[F(N_t,X_t,A_t)-\Xi\left(e_t\right)-W_tN_t\Big]\Bigg|\left(A_0,W_0\right)\right)$$

subject to

$$X_t \leq \underline{X}_t,$$

$$N_t = (1-d)N_{t-1}+e_t,$$

$$N_{-1} \text{ given.}$$

Let Q_t be the current-value multiplier associated with the first constraint and v_t be the multiplier for the second. Then an optimal allocation must satisfy the following conditions:

$$N_t \quad : \quad \frac{\partial F(N_t,\underline{X}_t,A_t)}{\partial N_t} = W_t + \left[v_t - \mathbf{E}_t\beta(1-d)v_{t+1}\right] \tag{16}$$

$$X_t \quad : \quad \frac{\partial F(N_t,\underline{X}_t,A_t)}{\partial X_t} = Q_t \tag{17}$$

$$e_t \quad : \quad v_t = \frac{\partial \Xi_t}{\partial e_t}. \tag{18}$$

We impose an exact equivalence of our model of discrete individual choice to the partial adjustment model by assuming (a) that $\underline{X}_t = \underline{X}$, $t = 0, 1\ldots,$ (b) that $F\left(N, \underline{X}, A\right) = f\left(N, A\right)$ for all N, A and (c) that total adjustment costs are quadratic as in (15). Under these assumptions (17) determines the rental price Q_t at given stock $\underline{X}_t$, and (16) and (18) are equivalent to (5) and (6). Since (3) describes the evolution of aggregate employment in both models, it follows that our discrete choice model is identical in its aggregate quantities to the traditional partial adjustment model. Therefore, if, as before, we now impose that $f(N, A)$ takes the form indicated in (8), then

$$\frac{\partial F(N_t, \underline{X}, A_t)}{\partial N_t} \approx f_n - f_{nn} N_t + f_{na} A_t$$

and our discrete individual choice model is described, at the industry-level, by (1) with N^* given by (10) and κ by (11).

4.2 Choices of individual production units

We now describe how to decentralize the planning solution above as a competitive equilibrium. In the setting that we have described, an individual production unit makes two sets of decisions. First, if it has already entered, then it chooses its factor demands so as to maximize its profits. Second, if it has not already entered, it decides whether to do so by comparing the present discount value of profits to its entry cost.

We assume that factor inputs are traded in competitive factor markets and begin by establishing that the optimal allocations satisfy the factor demands of active production units. For any production unit, profits are

$$\pi\left(n_t(z), x_t(z); A_t, W_t\right) = F(n_t(z), x_t(z), A_t) - W_t n_t(z) - Q_t x_t(z),$$

where the multiplier, Q_t, is now interpreted as the market price of the fixed factor. Given the unit capacity constraint on employment at any location, a production unit will hire labor to capacity and purchase $\underline{x}_t = \underline{X}_t / N_t$ units of the specific factor if $\frac{\partial F(1, \underline{x}_t, A_t)}{\partial n_t} > W_t$ and $\frac{\partial F(1, \underline{x}_t, A_t)}{\partial x_t} = Q_t$. Since the optimal allocation satisfies (16) and (17), these conditions are satisfied, and the choice is individually rational for an active production unit, provided $v_t - \beta(1-d)\mathbf{E}_t v_{t+1} > 0$. For every active production unit z, $n_t(z) = 1$, $x_t(z) = \underline{x}_t$ and profits are $\pi\left(1, \underline{x}_t, A_t, W_t\right)$.

Next we establish that the rate of entry in the optimal allocation is consistent with individual choice. Suppose that a particular potential production location faces a fixed cost of ξ_t. If this cost is smaller than the expected present discounted value of profits that will arise from entering, i.e.,

$$\xi_t \leq \mathbf{E}\left(\sum_{j=0}^{\infty}\Big(\beta(1-d)\Big)^j\Big[\pi\left(1,\underline{x}_{t+j};A_{t+j},W_{t+j}\right)\Big]\Bigg|\left(A_t,W_t\right)\right), \tag{19}$$

then the unit will benefit from paying the fixed cost and entering into productive activity. In this expression, future profits are discounted by $\beta(1-d)$ since the unit understands that the probability of remaining active without paying another fixed cost is $(1-d)$.

The present discounted value of profits can be expressed in two other ways. First, using the properties of the constant returns-to-scale production and efficient demand for the specific factor, (17), it follows that $F(1,\underline{x}_t,A_t)-W_t-Q_t\underline{x}_t=\frac{\partial F(1,\underline{x}_t,A_t)}{\partial n_t}-W_t$. Consequently, the right-hand side of (19) may be rewritten as

$$E\left(\sum_{j=0}^{\infty}\Big(\beta(1-d)\Big)^j\left[\left(\frac{\partial F(1,\underline{x}_{t+j},A_{t+j})}{\partial n_{t+j}}-W_{t+j}\right)\right]\Bigg|\left(A_t,W_t\right)\right),$$

which captures the idea that these profits are related to excess of the marginal product over the wage rate.[8] Second, in an optimal allocation, this present discounted value is just the multiplier v_t, as an implication of efficient adjustment of N_t.[9] Thus, a production unit will enter if its fixed cost is less than the present discounted value of profits, which are in turn summarized by a Lagrangian multiplier that may be interpreted as the market value of an active production unit: $\xi_t \leq v_t$. It then follows that, in competitive equilibrium,

[8] This gap arises because of the capacity constraint on labor. If the capacity constraint were at χ, then the entry decision would be based on the present discounted value of $\left(\frac{\partial F(\chi,\underline{x}_t,a_t)}{\partial n_t}-W_t\right)\chi$.

[9] This is derived using forward-iteration on (16) and the endpoint condition $\lim_{j\to\infty}\left(\beta\left(1-d\right)\right)^j v_{t+j+1}=0$.

the total number of entrants will be determined by the condition $v_t = \frac{\partial \Xi_t}{\partial e_t} \equiv \xi(e_t)$, where by $\xi(e_t)$ we mean the fixed cost faced by the marginal production unit if a fraction e_t of the production units is entering, $G^{-1}(e_t)$. This is exactly the condition determining the optimal adjustment rate in (18).

Collecting the results of this section, we conclude that, at the industry level, our stochastic adjustment model is observationally equivalent to the standard model given suitable (common) choices of parameters. It nevertheless has very different implications for establishment-level behavior, as emphasized below.

4.3 Matching the stylized facts

The stochastic adjustment model developed in this section matches the first two sets of the stylized facts that we highlighted in section 3. First, it captures the Hamermesh facts: (i) at the level of individual production units, there is discrete and occasional adjustment; (ii) heterogeneity in the circumstances of individual units, arising from differential fixed costs of adjustment, lead these actors to adjust at different times; and (iii) the industry adjustment process is well-approximated by the standard partial adjustment model. Second, it captures the first of the Caballero, Engel and Haltiwanger (1997) findings: the probability of adjustment for an individual production unit varies with aggregate conditions: in fact, a higher (lower) degree of aggregate adjustment in employment occurs only if there is a higher (lower) probability of individual adjustment.

5 Generalized partial adjustment

A number of recent theoretical and empirical studies – notably those of Caballero and Engel (1993) and Caballero, Engel and Haltiwanger (1997) – have argued for a richer vision of the adjustment process that can generate the remaining stylized facts discussed above. The framework of this section exemplifies such a model. In particular, our model delivers the implication that an individual production unit's probability of adjustment depends on a measure of the 'gap' between its current employment and a notion of desired employment, (fact 4). Further, it can produce substantial responses of employment to aggregate shocks without relying on any shifts in the distribution of adjustment cost, (fact 5). At the same time, the framework can be readily incorporated into a general equilibrium model, so that the relationship between adjustment hazards and macroeconomic conditions can be studied. In this section, we provide the essential elements of this framework and discuss its partial adjustment properties.

To achieve consistency with fact 4, we abandon the assumption that all active production units operate with the same level of employment. Instead, we now assume a large and fixed number of units, each making discrete choices about their employment adjustment, rather than the activation of production, over time. Production at the plant level is constant returns in labor and a fixed input, which we normalize to 1, $f(n_t, 1, A_t)$.[10] The presence of the fixed input allows determination of employment choice (relaxing the labor capacity constraint) at the production unit. Any unit that does not adjust its workforce sees it decay at rate d,

$$n_t = (1 - d)n_{t-1} + e_t, \tag{20}$$

where e_t is the number of hires.

We begin by assuming that the opportunity to adjust employment arrives exogenously according to a probabilistic mechanism specified below. (This assumption will be relaxed

[10] Here we have shifted the decreasing returns-to-scale from the industry-level to the level of the production unit, which can be rationalized by assuming that the fixed input is fixed across units.

in section 5.1.) To capture the observation that a production unit may have a greater likelihood of adjusting employment when there has been a longer interval since its last adjustment, we allow the probability to depend on the length of time since the unit last changed employment, which we index by j. That is, if a production unit has not adjusted its employment for $j-1$ periods, then the conditional probability of its being allowed to adjust its employment in the j^{th} period is α_j. For now, we assume that these adjustment probabilities depend only upon time since last adjustment and are a fixed vector $\boldsymbol{\alpha} = [\alpha_1, \alpha_2, \cdots \alpha_{J-1}, 1]$ over time. We further assume that $\alpha_{j-1} < \alpha_j$ for all $j = 1, 2, ..., J$, where J represents the maximum interval before a production unit will be allowed to adjust its employment with probability 1: $\alpha_J = 1$.

Let $n_{j,t}$ represent the current labor stock of a production unit that last adjusted its employment j periods ago. We use the notation $V_j(n_{j,t}, A_t, W_t)$ to denote the value of a production unit that last adjusted j periods ago, entering the current period with a workforce of $n_{j,t}$, that is not currently able to adjust its employment, and use $V_0(A_t, W_t)$ to denote the value of a production unit currently able to adjust. For a unit that is currently readjusting its stock of labor,

$$V_0(A_t, W_t) = \max_{n_t} \Big(f(n_t, A_t) - W_t n_t + \beta \mathbf{E} \Big[\alpha_1 V_0(A_{t+1}, W_{t+1}) + (1-\alpha_1) V_1((1-d) n_t, A_{t+1}, W_{t+1}) \mid A_t, W_t \Big] \Big), \tag{21}$$

where n_t is freely chosen. The right-hand side of the Bellman equation involves three expressions. First, there is the flow of current profit. Second, there is the discounted value of being a unit that adjusts next period, which occurs with probability α_1. Third, there is the value of being a unit that does not adjust next period, an outcome that occurs with probability $(1-\alpha_1)$.

For units not currently able to adjust their workforce, there are no decisions in this simple model, although there would be in more elaborate settings that allowed for adjustments on other margins, such as in hours-per-worker. Their value functions obey the functional equation

$$
\begin{aligned}
V_j(n_{j,t}, A_t, W_t) = f(n_{j,t}, A_t) - W_t n_{j,t} + \beta \mathbf{E}\Big[& \alpha_{j+1} V_0(A_{t+1}, W_{t+1}) \\
& + (1-\alpha_{j+1})\, V_{j+1}((1-d)\, n_{j,t}, A_{t+1}, W_{t+1}) \mid A_t, W_t \Big].
\end{aligned} \tag{22}
$$

Note that, for non-adjusting production units, labor evolves according to $n_{j,t} = (1-d)n_{j-1,t-1}$, reflecting the consequences of worker departures from the production unit.

The adjusting production units choose employment so as to maximize the right-hand side of (21), which results in an efficiency condition of the following form:

$$
D_1 f(n_t, A_t) - W_t + \beta \mathbf{E}\Big[(1-\alpha_1)\,(1-d)\, D_1 V_1(n_{1,t+1}, A_{t+1}, W_{t+1}) |\, A_t, W_t \Big] = 0.
$$

A notable feature of this condition is that it indicates that the optimal employment decision on the part of the adjusting production unit is independent of the length of time since it last adjusted and the size of its workforce at the start of the period, since neither j nor $n_{j,t}$ enters into the efficiency condition. This justifies our writing V_0 above in the restricted form that omits these factors. Working with the value function (22) above, we can determine the marginal value of additional workers:

$$
\begin{aligned}
D_1 V_j(n_{j,t}, A_t, W_t) &= D_1 f(n_{j,t}, A_t) - W_t \\
&\quad + \beta \mathbf{E}\Big[(1-\alpha_{j+1})\,(1-d)\, D_1 V_{j+1}(n_{j+1,t+1}, A_{t+1}, W_{t+1}) |\, A_t, W_t \Big].
\end{aligned}
$$

These derivatives may be used iteratively to simplify the efficiency condition and derive an alternative implicit expression for optimal choice of workforce chosen by an adjusting production unit. In particular, n_t^* solves

$$D_1 f(n_t, A_t) + \mathbf{E}\sum_{j=1}^{J-1}\Big[[\beta(1-d)]^j \varphi_j \Big(D_1 f((1-d)^j n_t, A_{t+j}) - W_{t+j}\Big)\Big| A_t, W_t\Big] = W_t, \quad (23)$$

where φ_j gives the adjusting unit's probability of remaining in the nonadjustment state for j consecutive periods:

$$\varphi_j \equiv \prod_{k=1}^{j}(1-\alpha_k),\ j = 1, \ldots, J-1. \quad (24)$$

It is now feasible to illustrate several features of our generalized partial adjustment model. Our first result establishes a property of the traditional partial adjustment model that survives our generalization. Specifically, equation (23) indicates that, when individual labor adjustments are discrete and occasional, an adjusting unit's labor demand, n_t, is a function of its expectation of future wages and productivities. Recall this property characterized target employment for the representative firm, equation (10), in the traditional partial adjustment model. However, in our generalized partial adjustment model, future adjustment probabilities also appear.

Our next result is easiest to illustrate when we set wages and productivities constant over time and examine the model's steady state. In such a deterministic setting with unchanging productivities and wages, the desired employment level is a constant $n^*(\boldsymbol{\alpha}, A, W)$. In this environment, our assumption that adjustment probabilities rise with time since last adjustment is equivalent to assuming that they are rising in the distance of current employment from the desired employment level. That is, the ordering of time since adjustment across production units is equivalent to the ordering by employment gap, since

$$n_j - n^*(\boldsymbol{\alpha}, A, W) = [(1-d)^j - 1]n^*(\boldsymbol{\alpha}, A, W).$$

This allows us to establish our second result, a *hedging* property in our generalized partial adjustment model that arises because of forecasted future labor force departures. Suppress expectations and time subscripts in (23), and note that, since the production function is concave in employment, the optimal employment choice will exceed the static optimum. Specifically, let n^s represent the static optimum that would be chosen if the unit could adjust its employment in every period with certainty; this static optimum satisfies $D_1 f(n^s, A) - W = 0$. Since $D_{11} f < 0$, it follows that for $j = 0, \ldots, J-2$,

$$D_1 f((1-d)^j n, A) - W < D_1 f((1-d)^{j+1} n, A) - W.$$

Hence the summation in (23) evaluated at $n = n^s$ will be positive, proving that the static optimum cannot be the dynamic optimum. Moreover, as both this sum and its preceding expression, $D_1 f(n, A) - W$, are decreasing in n, it follows that the dynamic optimum, n^*, exceeds n^s. This is the hedging motive that raises employment above the static optimum. Production units hire more labor than they currently need in an effort to offset the probability that they may be unable to hire in the immediate future. Further, it is clear that n^* will be larger the higher is this probability of future nonadjustment; for instance, given d and $\alpha_2 \cdots \alpha_{j-1}$, a lower α_1, implying a reduced probability of adjustment in the first period after an adjustment, implies higher values for $\varphi_1, \ldots, \varphi_{J-1}$ and thus a higher value for the summation for any n. Consequently, the higher is the probability of being unable to restock employment, the stronger is the hedging motive.

To further illustrate the hedging motive, we examine the case of a Cobb-Douglas production function, $y = An^\gamma$. Equation (23) may be explicitly solved for the optimal labor demand for adjusting production units as

$$n^* = \left(\frac{\gamma A}{W}\right)^{\frac{1}{1-\gamma}} \left[\frac{1 + \sum_{j=1}^{J-1} b^j \varphi_j (1-d)^{\gamma j}}{1 + \sum_{j=1}^{J-1} b^j \varphi_j (1-d)^j}\right]^{\frac{1}{1-\gamma}}.$$

The first term in this expression, $[\frac{\gamma A}{w}]^{\frac{1}{1-\gamma}}$, is the standard static demand for labor that arises with a Cobb-Douglas production function. The presence of the second term may be exposited as follows. Consider a production unit that has adopted the statically optimal level of employment. Looking forward one period, the production unit knows that a fraction $(1-d) < 1$ of its workforce will be retained, which has the effect of lowering its wage bill and reducing its stock of workers. However, since the marginal product will increase as these workers depart, the future marginal product will exceed the wage rate. It is therefore optimal to raise n^* relative to $[\frac{\gamma A}{W}]^{(\frac{1}{1-\gamma})}$.

5.1 Endogenizing adjustment

We now endogenize the timing of individual production units' adjustment by introducing fixed costs of adjustment that are stochastic across production units, an approach adopted by Caballero and Engel (1999) in their study of manufacturing investment. Within each date, any individual production unit faces a random cost ξ that it must pay in order to adjust its employment. As in the entry model of section 4, this cost is drawn from a time-invariant distribution over $[0, B]$ that is summarized by the CDF $G(\xi)$ and associated PDF $g(\xi)$.

At the start of each date t, any establishment may be identified as a member of a particular time-since-adjustment group, j, where j indicates the numbers of periods that have elapsed since the last active employment adjustment. Given its cost draw of ξ, and given the aggregate state, such an establishment will adjust its employment if the fixed cost does not exceed the value of the adjustment, that is, if $V_0(A_t, W_t) - V_j(n_{jt}, A_t, W_t) \geq \xi$. Because there is a large number of production units within each different time-since-adjustment group, each group is characterized by a marginal plant that finds it just worthwhile to adjust. This marginal plant is associated with a cost $\overline{\xi}_{jt}$ such that

$$\overline{\xi}_{jt} = V_0(A_t, W_t) - V_j(n_{jt}, A_t, W_t). \tag{25}$$

All production units in the j^{th} time-since-adjustment group with adjustment costs at or below the threshold in (25) will choose to adjust. As a result, the fraction of plants adjusting out of any particular group j, $j = 1, \ldots, J-1$, is given by

$$\alpha_{jt} = G(\overline{\xi}_{jt}). \tag{26}$$

From (25), note that these adjustment fractions are functions of the plant-level state vector, (n_{jt}, A_t, W_t). We assume that the stochastic processes for productivity and wages are such that, given the function f and the discount factor β,

$$B < V_0(A_t, W_t) - V_J(n_{Jt}, A_t, W_t)$$

for all values of the vector (n_{Jt}, A_t, W_t). This assumption, which follows quite naturally from $B < \infty$, given bounded processes A_t and W_t, assures us that $\alpha_J = 1$.

Having described the determination of endogenous adjustment probabilities, we must restate the plant's optimization problems to introduce adjustment costs and time-varying adjustment probabilities determined by (25) - (26). With state-dependent probability $\alpha_{j+1,t+1}$, a production unit entering period $t+1$ in group $j+1$ will adjust at that date. The counterpart to (21), the value of a plant that is currently adjusting its labor, is

$$\begin{aligned} V_0(A_t, W_t) = \max_{n_t} \Big(f(n_t, A_t) - W_t n_t + \beta \mathbf{E}\Big[\alpha_{1,t+1} \Big(V_0(A_{t+1}, W_{t+1}) - \xi_{1,t+1} \Big) \\ + (1 - \alpha_{1,t+1})\, V_1((1-d)\, n_t, A_{t+1}, W_{t+1}) \mid A_t, W_t \Big] \Big), \end{aligned} \tag{27}$$

where $\xi_{1,t+1}$ reflects the expected fixed cost that the plant will pay at date $t+1$, conditional on its undertaking an employment adjustment. Similarly, the value of a non-adjusting plant that last adjusted j periods ago, the counterpart to equation 22, is

$$V_j(n_{j,t}, A_t, W_t) = f(n_{j.t}, A_t) - W_t n_{j,t} + \beta \mathbf{E}\Big[\alpha_{j+1,t+1}\Big(V_0(A_{t+1}, W_{t+1}) - \xi_{j+1,t+1}\Big) \qquad (28)$$

$$+ (1 - \alpha_{j+1,t+1}) V_{j+1}((1-d)\, n_{j,t}, A_{t+1}, W_{t+1}) \mid A_t, W_t\Big].$$

Adjusting plants exit the j^{th} group for the adjustment group and choose an optimal employment level n_t^* (or $n_{0,t}$) satisfying the marginal profit condition in (29), which generalizes (23) to reflect optimal adjustment probabilities:

$$D_1 f(n_t, A_t) - W_t + \mathbf{E} \sum_{j=1}^{J-1} \Big[[\beta(1-d)]^j \varphi_{j,t+j} \{ D_1 f((1-d)^j n_t, A_{t+j}) - W_{t+j} \Big| A_t, W_t \Big] = 0. \qquad (29)$$

Here, as in (24), $\varphi_{j,t+j}$ is the probability the unit will make no further adjustment in the next j periods. That is, for $j = 1, ..., J-1$,

$$\varphi_{j,t+j} \equiv \prod_{k=1}^{j} \Big(1 - \alpha_{k,t+k}\Big) = \prod_{k=1}^{j} \Big(1 - G(\overline{\xi}_{k,t+k})\Big). \qquad (30)$$

5.2 Partial adjustment of market labor demand

The probabilistic approach to microeconomic employment adjustment that we have constructed is consistent with the empirical evidence on rising employment adjustment hazards. Moreover, the framework allows us to aggregate individual plants' labor demand and derive a simple expression for market labor demand. Since the economy is populated by a large number of production units, we can describe the distribution of plants in any date t using the vector $\theta_t = [\theta_{1,t}, ..., \theta_{J,t}]$, with each $\theta_{j,t}$ representing the fraction of units that begin the period having last adjusted j periods prior to the current date.[11] Letting

[11]More precisely, the distribution at the start of any date t is completely summarized by the vector θ_t together with a vector of previous target employment levels $[n_{t-1}^*, ..., n_{t-J}^*]$ from which the support is

$\omega_{0,t}$ denote total adjusting units in any date t, the elements of this vector are as follow.[12]

$$\theta_{1,t} = \omega_{0,t-1} \tag{31}$$

$$\theta_{j,t} = (1 - \alpha_{j-1,t-1})\,\theta_{j-1,t-1} \qquad \text{for} \quad j = 2, ..., J. \tag{32}$$

Market labor demand may then be represented as a moving average of the employment actions of production units, with lag weights determined by adjustment fractions across time-since-adjustment groups:

$$N_t = n_t^* \sum_{j=1}^{J} \theta_{j,t}\alpha_{j,t} + \sum_{j=1}^{J-1} \theta_{j,t}\,(1 - \alpha_{j,t})\, d^j n_{t-j}^*. \tag{33}$$

This is the third result of our generalized partial adjustment model. The market's dynamic demand for labor describes aggregate employment as a weighted average of past target employment as in the traditional partial adjustment model (2). Consequently, while the underlying production unit level demands are adjusted discretely and occasionally, the market demands vary smoothly in every time period. Given that each target employment, n_{t-j}^*, $j = 1, \ldots, J-1$, involves expectations of future wages and productivities, so does market labor demand.

While equation (33) shows that our generalized partial adjustment model has a representation similar to the traditional partial adjustment model, there are important differences that eliminate exact aggregate equivalence. In particular, the lag weights here vary

trivially retrieved.

[12] Given a fixed measure of production units, the fraction adjusting in date t may be expressed as $\omega_{0,t} = 1 - \sum_{j=1}^{J-1} \omega_{0,t-j}\varphi_{j,t}$, where each $\varphi_{j,t}$ is as defined in (30), with the appropriate date change. Thus, in the stationary distribution, the overall adjustment rate is $\omega_0 = \left(1 + \sum_{j=1}^{J-1} \varphi_j\right)^{-1}$.

over time, because they are composite functions of the adjustment rates α_j, which themselves are functions of plant and aggregate state variables, as consistent with stylized fact 3. Thus, in contrast to the traditional partial adjustment model, our economywide rate of adjustment responds to changes in aggregate conditions, including changes in economic policy.

5.3 A planning representation

The generalized partial adjustment model described above may be derived as the solution to a single dynamic optimization problem. We briefly outline this reformulation to illustrate the tractability of the approach and thus its suitability for applications.[13] The aggregate representation consolidates the ownership of all plants, differentiated by their time since last adjustment, $j = 1, \ldots, J$, into a single entity, a representative firm. Using the notation $\boldsymbol{\theta}_t \equiv [\theta_{1t}, ..., \theta_{Jt}]$, $\mathbf{n}_t \equiv [n_{1t}, ..., n_{J-1,t}]$, and $\boldsymbol{\alpha}_t \equiv [\alpha_{1t}, ..., \alpha_{Jt}]$ to describe the economywide distribution of plants, employment, and adjustment fractions across groups, the flow return for the firm is the total of output less wages and adjustment costs,

$$\Pi\left(\boldsymbol{\theta}_t, \mathbf{n}_t, \boldsymbol{\alpha}_t; A_t, W_t\right) = \sum_{j=1}^{J} \alpha_{jt}\theta_{j,t}\Big(f(n_{0,t}, A_t) - W_t n_{0,t}\Big) - \sum_{j=1}^{J} \theta_{jt}\Xi(\alpha_{jt}) \qquad (34)$$
$$+ \sum_{j=1}^{J-1}(1-\alpha_{jt})\theta_{j,t}\Big(f(n_{jt}, A_t) - W_t n_{jt}\Big),$$

where $\Xi(\alpha) = \int_0^{\xi(\alpha)} xg(x)dx$ is the total volume of costs averaged across plants in group j if a fraction α adjusts, and $\xi(\alpha)$ is the value of ξ such that $\alpha = G(\xi)$. Given the current distribution of plants over time-since-last-adjustment groups, and given the sequence

[13] Note that, in contrast to the ordering of exposition in section 4, here we have chosen to begin our discussion with a description of decentralized actions and now follow with a planning representation. The reverse ordering would have been equally straightforward, which emphasizes the flexibility of the approach. The representation is selected according to its convenience for a particular application or solution method.

of wages and productivities, the representative firm chooses fractions of plants adjusting $(\alpha_{jt})_{j=1}^{J}$ and optimal employment at those that are adjusting their workers, $n_{0,t}$. The planning problem is then

$$V(\boldsymbol{\theta}_t, \mathbf{n}_t; A_t, W_t) = \max_{\boldsymbol{\alpha}_t, \boldsymbol{\theta}_{t+1}, n_{0t}} \Pi(\boldsymbol{\theta}_t, \mathbf{n}_t, \boldsymbol{\alpha}_t; A_t, W_t) \tag{35}$$
$$+\beta \mathbf{E}\Big[V(\boldsymbol{\theta}_{t+1}, \mathbf{n}_{t+1}; A_{t+1}, W_{t+1}) \mid A_t, W_t\Big]$$

subject to

$$n_{j,t} = (1-d)n_{j,t-1},\ j = 1, \dots, J-1, \tag{36}$$

$$(30) - (32) \text{ and } (34).$$

Let v_{0t} be the multiplier associated with (31), and v_{jt} denote the multipliers associated with (32). Efficiency with respect to the choice of α_{jt} requires that the solution to this problem satisfy

$$\xi(\alpha_{jt}) = v_{0t} - v_{jt},$$

so that it is just worthwhile to relocate the marginal plant with cost ξ_{jt} into the adjustment group, and plants with costs greater than this threshold are not adjusted. This determines $\alpha_{jt},\ j = 1, \dots, J$, and is equivalent to (25) provided the multipliers v_{jt} attain the same value as before. That this is the case may be seen from the efficiency condition with respect to $\theta_{jt},\ j = 1, \dots, J$, which implies that the value associated with a plant with employment level n_{jt} satisfies

$$v_{jt} = f(n_{jt}, A_t) - W_t n_{jt} + \beta \mathbf{E}\Big[\Big(\alpha_{j+1,t+1} v_{0,t+1} - \Xi(\alpha_{j+1,t+1})$$
$$+ (1-\alpha_{j+1,t+1})\, v_{j+1,t+1}\Big) \mid A_t, W_t\Big]$$

and, since full adjustment will take place by group J,

$$v_{J-1,t} = f(n_{J-1,t}, A_t) - W_t n_{J-1,t} + \beta \mathbf{E}\Big[\Big(v_{0,t+1} - \Xi(1)\Big) \mid A_t, W_t\Big].$$

These expressions are equivalent to the plant Bellman equations of section 5.1 since the expected adjustment cost conditional on adjustment in (28), $\mathbf{E}\left(\alpha_{j+1,t+1}\xi_{j+1,t+1}\right) = \Xi\left(\alpha_{j+1,t+1}\right)$, the average cost paid by adjusting plants, by definition of $\Xi\left(\cdot\right)$. Finally, the efficiency condition with respect to the choice of n_{0t} may be expressed as (29). Therefore the solution to the firm's dynamic profit maximization problem, given the aggregate state $(\boldsymbol{\theta}_t, \mathbf{n}_t; A_t, W_t)$, is the same as in the decentralized model of the previous section.

5.4 Numerical examples: The five stylized facts

We use a series of numerical examples to illustrate several interesting properties of the generalized partial adjustment model developed above, and to contrast its dynamics to those of the traditional partial adjustment model. We begin with an examination of the model assuming that prices, wages and interest rates are exogenously fixed, as is commonly the case in analyses using the traditional partial adjustment model. Our examination involves functional forms and parameter values that are standard. Specifically, we assume that production at the plant level is described by a Cobb-Douglas production function $f(n, A) = An^{\nu}$ where $\nu = 0.66$. Total factor productivity has a mean of 1 and follows a first-order autoregressive process with a one-period autocorrelation of 0.9225, roughly consistent with the annual properties of the Solow Residual. The firm's discount factor is $\beta = 0.939$, which corresponds to an annual interest rate of 0.065. These values will be familiar to quantitative researchers; see, for example, King and Rebelo (1999) or Thomas (2002).

The remaining parameter values are chosen arbitrarily; however, extensive sensitivity analysis has confirmed that the properties of the model we have developed are not qualitatively sensitive to variation in these parameters. First, we assume that the distribution of adjustment costs is uniform with an upper support of 0.008. This yields a distribution of employment across plants that is suitable for illustrating the generalized partial adjustment model's properties. Next, for the partial adjustment model, we assume $B = 4$. This choice facilitates comparison, as it yields a dynamic response that is relatively close to our generalized partial adjustment model with adjustment rates held constant. Finally, we assume a separation rate of $d = 0.06$ and a wage rate of $W = 1.14$.

Before proceeding further, it is useful to note that we have developed a model that is designed to be consistent with stylized facts 1 and 5 of section 3. Specifically, due to fixed costs of adjustment, labor changes at the plant level are discrete and occasional in the model. Moreover, since the distribution of adjustment costs is assumed to be constant over time, it cannot be the source of aggregate fluctuations. Such fluctuations must arise through aggregate shocks as suggested by previous empirical work.

Our first figure showing the stationary distribution of plants illustrates the model's ability to reproduce stylized fact 4: adjustment probabilities depend on plants' gaps between actual and target employment. In figure 1, we see that adjustment fractions are an increasing function of the time since last adjustment, as the cost of non-adjustment rises with the level of disequilibrium, while the distribution of adjustment costs is identical across groups. Thus, in the second panel of the figure, the distribution function of firms across groups is necessarily downward sloping, given the law of motion for $\boldsymbol{\theta}$ in (32).

The next graph, figure 2, illustrates stylized fact 2; aggregate employment is characterized by smooth and gradual adjustment. Panels (a) and (b) show percentage deviations in market employment and output demand from their steady state values, in response to a persistent rise in aggregate productivity, for each of three models discussed above. PA corresponds to the traditional partial adjustment model of section 2, where staggered aggregate adjustment arises from the presence of quadratic adjustment costs, while TD represents the response for the generalized partial adjustment model with a fixed vector of time-dependent adjustment fractions. Finally, SD denotes the response in the generalized state-dependent partial adjustment model. In each of the two generalized partial adjustment models, fixed costs of adjustment dissuade some production units from responding at once to the rise in productivity. This protracts the aggregate response in employment and hence output, so that both TD and SD share the hump-shaped features that distinguish the partial adjustment model. This hump-shaped response in employment, most pronounced for the SD model, would be absent in a frictionless model of employment adjustment. There, without adjustment costs, the response of aggregate employment would be identical to the monotonic response of the auto-correlated productivity shock.

The time-dependent model, with an upward sloping but time-invariant adjustment hazard, matches the traditional partial adjustment model closely. Only at the earliest date of the response does the traditional model move more gradually than in TD, due to the rising marginal cost of aggregate employment changes. The size of this initial difference in employment response is nonetheless only about two-thirds of 1 percent. This is in part because plants in the time-dependent adjustment model are not permitted to alter the timing of their employment adjustments in response to shocks, so that all rises in aggregate employment must come from changes in intensive margin adjustment decisions. Moreover, the onset of diminishing returns at the level of the production unit restrains the rise in the employment levels chosen by current adjustors in TD.

While the state-dependent adjustment model shares similar qualitative features with the other staggered adjustment models, the ability of establishments to alter the timing of their employment adjustments at relatively low cost produces two potentially important changes in the market response. First, because aggregate employment is increased through changes in both intensive and extensive margin adjustment, SD produces a substantially larger rise in employment, and hence output, at the dates of highest productivity. It is precisely this 'time-varying elasticity' of aggregate employment demand with respect to aggregate shocks that distinguishes the SD model relative to the traditional model. The empirical work of Caballero and Engel (1992, 1993) finds that such properties are important in explaining the dynamics of aggregate employment demand during episodes involving unusually large shocks, such as the recession of 1974-1975 and the subsequent expansion. Second, the model has the ability to produce more complicated cyclical adjustment patterns; in each panel the SD response oscillates above and below the traditional model's response. As neither of these features in present when adjustment rates are held fixed, it is apparent that they arise due to large changes in adjustment timing at the micro-level.

Figure 3 verifies the importance of the time-varying plant distribution by displaying the SD responses in each of the two margins through which aggregate employment is raised. Panel (a) depicts percent changes in extensive margin adjustment through changes in the fraction of production units adjusting, $\omega_{0,t} = \sum_{j=1}^{J} \theta_{jt}\alpha_{jt}$, while panel (b) displays intensive margin changes through the employment levels chosen by current adjustors, n_{0t}. Given the persistent nature of the productivity shock, the rewards to early adjustment are expected to be large, thus raising the threshold costs above which adjustment is rejected within each time-since-adjustment group. As a result, adjustment fractions rise across groups, and the number of adjustors in the economy rises 25 percent above its steady state value. This illustrates that the third stylized fact is met by our generalized partial adjustment model: adjustment rates vary with aggregate conditions. Note that, in contrast to the large change in adjustment rates, the percent rise in target employment per adjusting unit is considerably smaller. Large increases in employment are not worthwhile given decreasing returns in establishment-level production. Thus, for this particular example, changes in the extensive margin, the number of adjusting plants, are more important than changes in the intensive margin, the level of employment chosen by such plants, in determining movements in aggregate employment. Furthermore, the extensive margin is responsible for the cyclical pattern seen in figure 2 for the aggregate quantity series.

Comparing panels (a) and (b) of figure 3, note that, while target employment monotonically declines with the decay of the shock, the number of adjustors oscillates in its return to steady state. The large rise in the number of adjustors at the impact of the shock results in a substantial shift in the distribution of production units away from higher time-since-adjustment groups and into group 1 starting the next period. Given the rising adjustment hazard, only a small fraction of these extra members find it worthwhile to again adjust their employment, so many of the initial surge in adjustors begin the subsequent date in group 2. In this way, the effects of early rises in adjustment rates filter out through subsequent distributions, reducing total adjustment toward trend, and then below it once a disproportionate fraction of the population finds its way into time-since-adjustment groups associated with low adjustment fractions. Eventually, the mass of early adjustors works its way sufficiently far out the distribution, where adjustment rates are relatively high, so that total adjustment returns above trend. This pattern is repeated in a dampened fashion until the distribution finally resettles.

Figure 4 aggregates the effects of changes in intensive margin versus extensive margin adjustment to provide a decomposition of the aggregate employment response into two underlying components: "n_j effects" associated with changes in employment levels across groups (due to changes in target employment chosen by adjustors) and "ω_j effects" arising from changes in the distribution of plants across these groups at the time of production, $\omega_{j,t} \equiv (1-\alpha_{jt})\theta_{jt}$, $j = 1, ..., J$, (due to changes in the fractions of units adjusting from each group). Specifically, at each date, the percentage deviation from steady state in aggregate employment is given by

$$\widehat{n}_t = \left[\sum_{j=0}^{J-1} \left(\frac{\omega_j n_j}{n} \right) \widehat{n}_{jt} \right] + \left[\sum_{j=0}^{J-1} \left(\frac{\omega_j n_j}{n} \right) \widehat{\omega}_{jt} \right],$$

where each $\left(\frac{\omega_j n_j}{n}\right)$ reflects the percentage contribution of the j^{th} group to aggregate employment in steady state, and each $\widehat{n}_{jt}$ and $\widehat{\omega}_{jt}$ represent percent deviations from trend in the group j employment and population levels, respectively, at the time of production in date t. At the onset of the shock, rises in employment associated with current adjustors, n_{0t}, contribute less than half of the percentage rise in the aggregate series. The remainder is due to a rise in the adjustment group, ω_0, associated with this high target and corresponding reductions in the populations of groups associated with lower employment levels, ω_j, $j = 1, \ldots, J$. In the following date, adjusting plants again select a high target employment level, and this is compounded by a rise in the employment held by members of group 1, a consequence of the high employment choice of the previous period. These effects of raised targets continue to feed through the distribution, raising the employment levels associated with each subsequent group, for a number of periods. As a result, the n_j component of aggregate employment exhibits the smooth humped shape associated with partial-adjustment. The aggregate series inherits this shape to an extent, but it is both more pronounced in its rise and less smooth in its return to trend, due to the ω_j effects arising from changes in membership across groups. High adjustment fractions amplify the aggregate response initially; however, by date 3, when the number of adjustors begins to fall below trend, an increasing fraction of production units operates with relatively low employment levels. This dampens the rise in the aggregate series, and speeds its initial rate of decline, relative to that of the n_j component. Further, just as the disruption in the population distribution produced oscillations in the total adjustors' series of figure 3, it also causes overshooting in the ω_j component's convergence and thereby generates the cyclical features evident in the aggregate series.

5.5 Extension to equilibrium

It is straightforward to embed the representative firm's problem from section 5.3 within a fully specified equilibrium model.[14] We briefly examine such an equilibrium version of our generalized partial adjustment model here. Assume that a representative household maximizes present discounted lifetime utility,

$$\mathbf{E}_0 \sum_{t=0}^{\infty} \beta^t U\left(C_t, N_t\right), \tag{37}$$

subject to $C_t = \Pi_t + W_t N_t$, where N_t is given by (33). Choosing $\boldsymbol{\alpha}_t$ and n_{0t} to now maximize expected lifetime utility, (37), rather than expected lifetime profits, (35), subject to (31) - (32), (34) and (36), we obtain equilibrium employment dynamics. To illustrate these, we assume $U\left(C, N\right) = \log\left(C - \chi \frac{N^{1+\gamma}}{1+\gamma}\right)$, where we choose $\chi = 2.55$ and $\gamma = 0.50$. This implies a steady state hours worked of 0.20 and a fairly high elasticity of labor supply. Higher values of γ would imply sharper differences between the equilibrium and fixed price models.

Figure 5 compares the response of the state-dependent generalized partial adjustment model in equilibrium to that occurring in the absence of price movements. Quantitatively, as might be expected, equilibrium price movements sharply dampen the response in employment, and hence output, to a persistent change in productivity. However, important

[14] The generalized adjustment model developed above has been used in several general equilibrium applications. Dotsey, King and Wolman (1999) study the dynamics of price adjustment, while Thomas (2002) and Khan and Thomas (2003) investigate investment dynamics. In this study, we use linear approximation methods in the tradition of Sargent (1978) to explore the general equilibrium connections, as do Dotsey, King and Wolman and Thomas. However, such models can be solved using alternative approximation methods, as in Khan and Thomas.

qualitative changes also arise, in both the extensive and intensive margins of employment adjustment. First, the previous nonmonotonicity in the fraction of units adjusting disappears in equilibrium. Second, the smooth mean reversion in target employment becomes less regular.

6 Concluding remarks

Using a time-invariant distribution of adjustment costs that are random across production units at a point in time, and over time for any unit, we have developed a new variety of partial adjustment model for labor demand. Our generalized partial adjustment model is consistent with 5 stylized facts: (1) employment adjustment at the establishment is discrete and occasional, (2) aggregate employment is smooth and gradual, (3) individual plants' probabilities of adjustment, their adjustment rates, vary over time in response to aggregate conditions, (4) these adjustment probabilities are functions of the difference between plants' actual and target employment and (5) movements in aggregate employment are largely driven by movements in aggregate factors, not by changes in plant-level factors.

The last stylized fact has motivated our focus on idiosyncratic uncertainty at the plant level that is transitory and our abstraction from additional sources of plant-specific heterogeneity. Existing empirical research suggests that such factors are of secondary importance in explaining movements in aggregate employment. A benefit to our abstraction is that we are able to develop a generalized (S, s) model of establishment-level labor adjustment that rationalizes existing empirical work that has heretofore assumed state-dependent adjustment hazards. Moreover, we have shown that our method allows convenient aggregation of the discrete adjustment actions of a heterogeneous distribution of production units into a smooth decision problem of a single representative firm.

Using our generalized partial adjustment model, we have analyzed the dynamics of employment under two alternative assumptions about the wage rate and interest rate, two prices that are central to an establishment's adjustment decision. We began by assuming that both prices were fixed, while productivity fluctuated exogenously. Next we considered a simple general equilibrium formulation in which these prices were endogenously determined, and hence varied with changes in productivity. The dynamics under these two formulations are quite different, but the differences are understandable consequences of variations in wages and interest rates. Previous research in this area has been conducted almost exclusively under the assumption of exogenous prices, given the complications presented by nontrivial heterogeneity in production. An important contribution of the current model lies in its ability to limit such complications, thereby allowing straightforward aggregation, and hence the natural extension to general equilibrium. We therefore view it as a tractable basis for future research into the dynamics of factor adjustment.

References

[1] Caballero, Ricardo J. and Eduardo M. R. A. Engel [1992], "Beyond the Partial-Adjustment Model," *American Economic Review* 82, 360-384.

[2] Caballero, Ricardo J. and Eduardo M. R. A. Engel [1993], "Microeconomic Adjustment Hazards and Aggregate Dynamics," *Quarterly Journal of Economics* 82, 359-383.

[3] Caballero, Ricardo J. and Eduardo M. R. A. Engel [1999], "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S, s) Approach," *Econometrica* 67, 783-826.

[4] Caballero, Ricardo J.; Eduardo M. R. A. Engel; and John C. Haltiwanger [1997], "Aggregate Employment Dynamics: Building From Microeconomic Evidence," *American Economic Review* 87, 115-137.

[5] Dotsey, Michael; Robert G. King; and Alexander L. Wolman, [1999], "State Dependent Pricing and the General Equilibrium Dynamics of Money and Output," *Quarterly Journal of Economics* 14, 655-690.

[6] Hamermesh, Daniel S. [1989], "Labor Demand and the Structure of Adjustment Costs," *American Economic Review* 79, 674-689.

[7] Hamermesh, Daniel S. and Gerard A. Pfann [1996], "Adjustment Costs in Factor Demand," *Journal of Economic Literature* 34, 1264-1292.

[8] Khan, Aubhik and Julia K. Thomas [2003], "Nonconvex factor adjustments in equilibrium business cycle models: do nonlinearities matter?" *Journal of Monetary Economics* 50, 331-360.

[9] King, Robert G. and Sergio T. Rebelo [1999], "Resuscitating Real Business Cycles," chapter 14 in M. Woodford and J. Taylor (eds.) *Handbook of Macroeconomics* IB Elsevier Science.

[10] Kollintzas, Tryphon E. [1985], "The Symmetric Linear Rational Expectations Model," *Econometrica* 53, 963-976.

[11] Lucas, Robert E. Jr. and Edward C. Prescott [1971], "Investment Under Uncertainty," *Econometrica* 39, 659-681.

[12] Mortensen, Dale T. [1973], "Generalized Costs of Adjustment and Dynamic Factor Demand Theory," *Econometrica* 41, 657-665.

[13] Parente, Stephen L. and Edward C. Prescott [1994] "Barriers to Technology Adoption and Development," *Journal of Political Economy* 102, 298-321.

[14] Sargent, Thomas J. [1978], "Estimation of Dynamic Labor Demand Schedules under Rational Expectations," *Journal of Political Economy* 86, 1009-1044.

[15] Thomas, Julia K. [2002], "Is Lumpy Investment Relevant for the Business Cycle?" *Journal of Political Economy* 110, 508-534.

Figure 1a: Steady State Adjustment Fractions

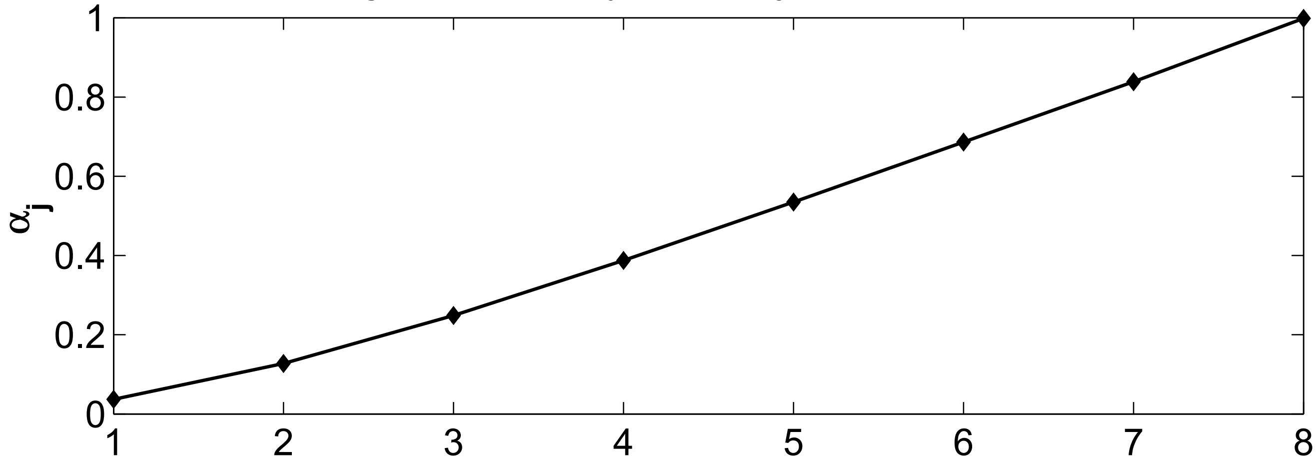

Figure 1b: Steady State Distribution

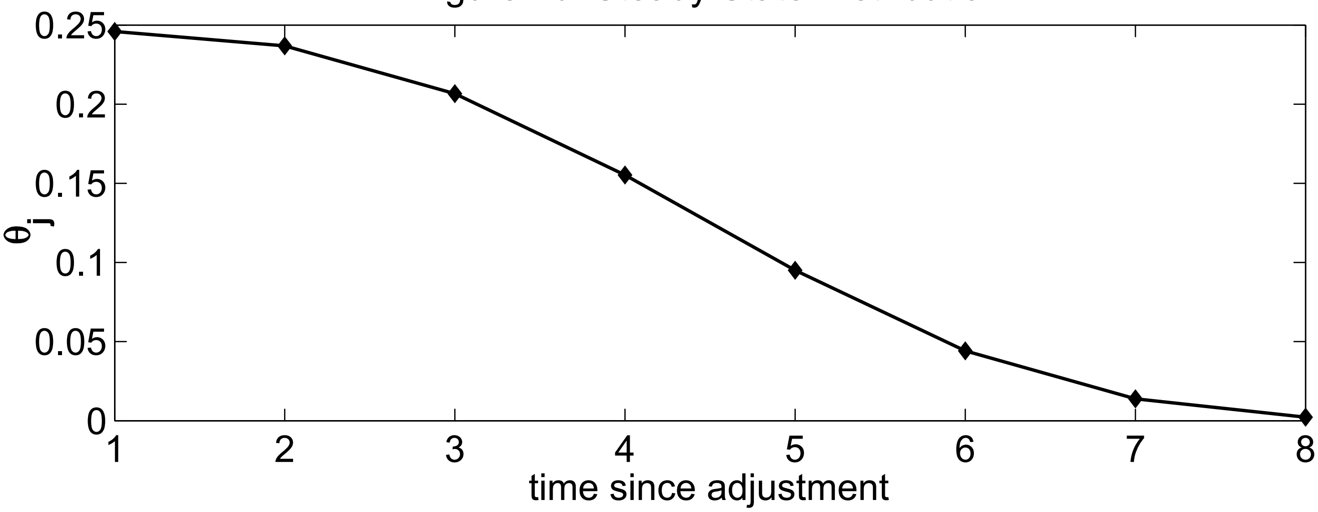


Figure 2a: Market Employment
percent deviation
PA
TD
SD
0
1
2
3
0
2
4
6
8
10

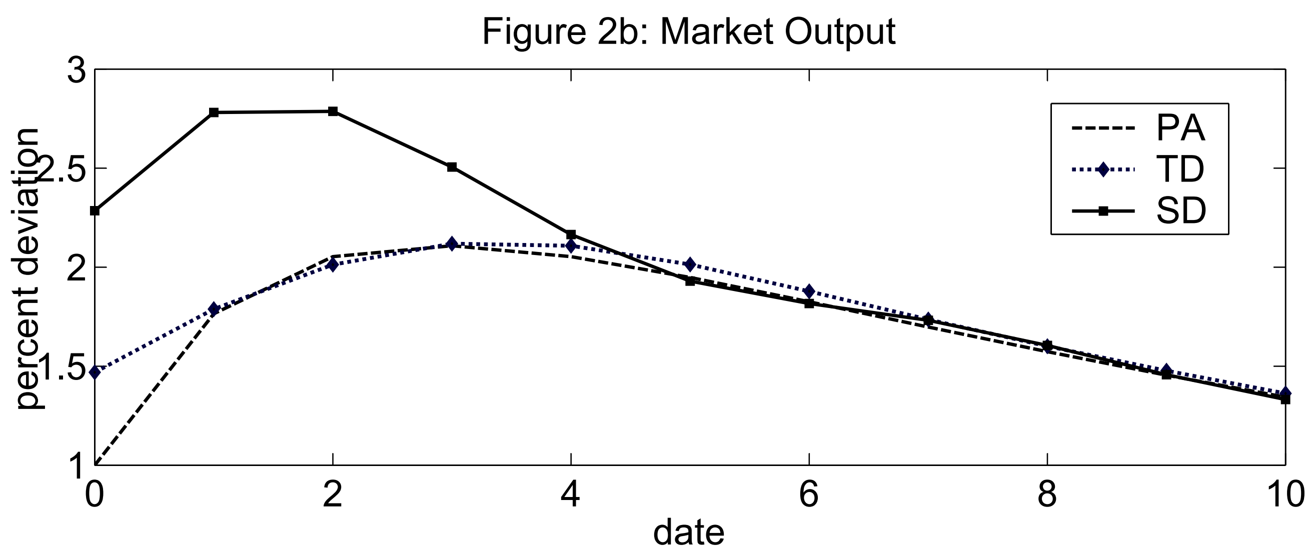
Figure 2b: Market Output
percent deviation
PA
TD
SD
1
1.5
2
2.5
3
0
2
4
6
8
10
date

Figure 3a: Total Adjustors

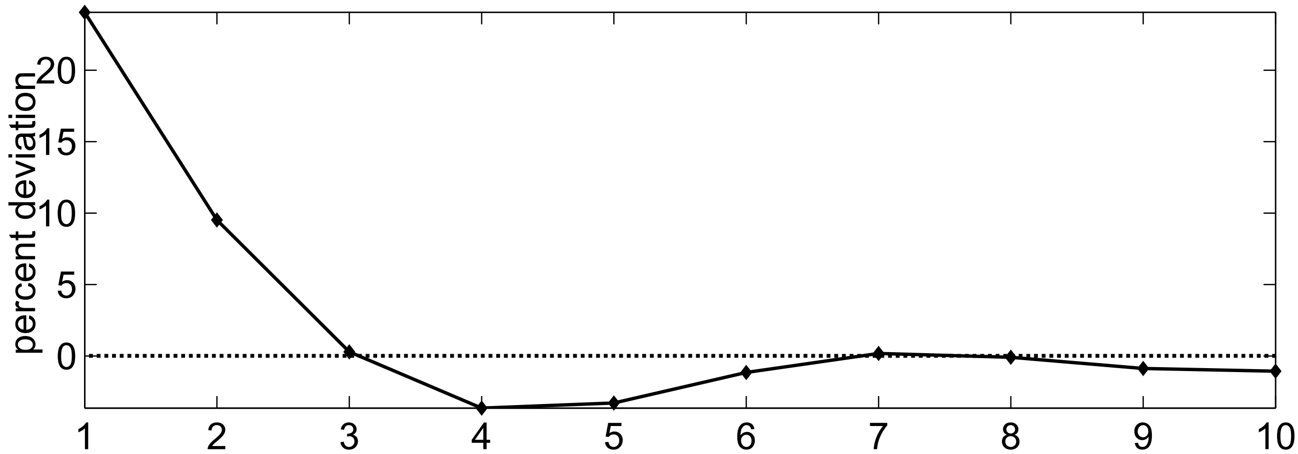

Figure 3b: Target Employment

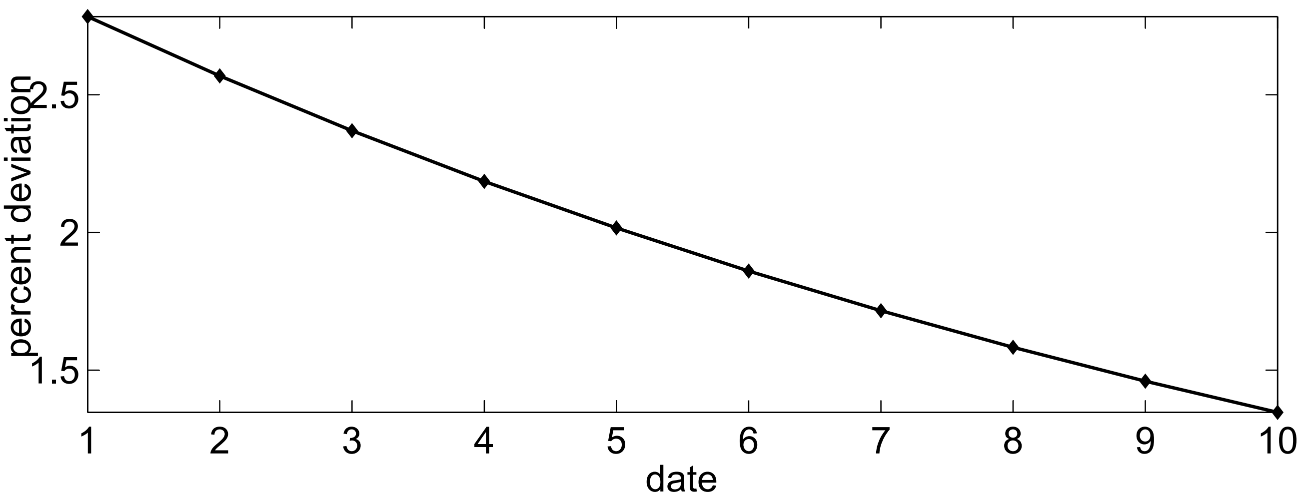

Figure 4: Decomposition of Aggregate Employment

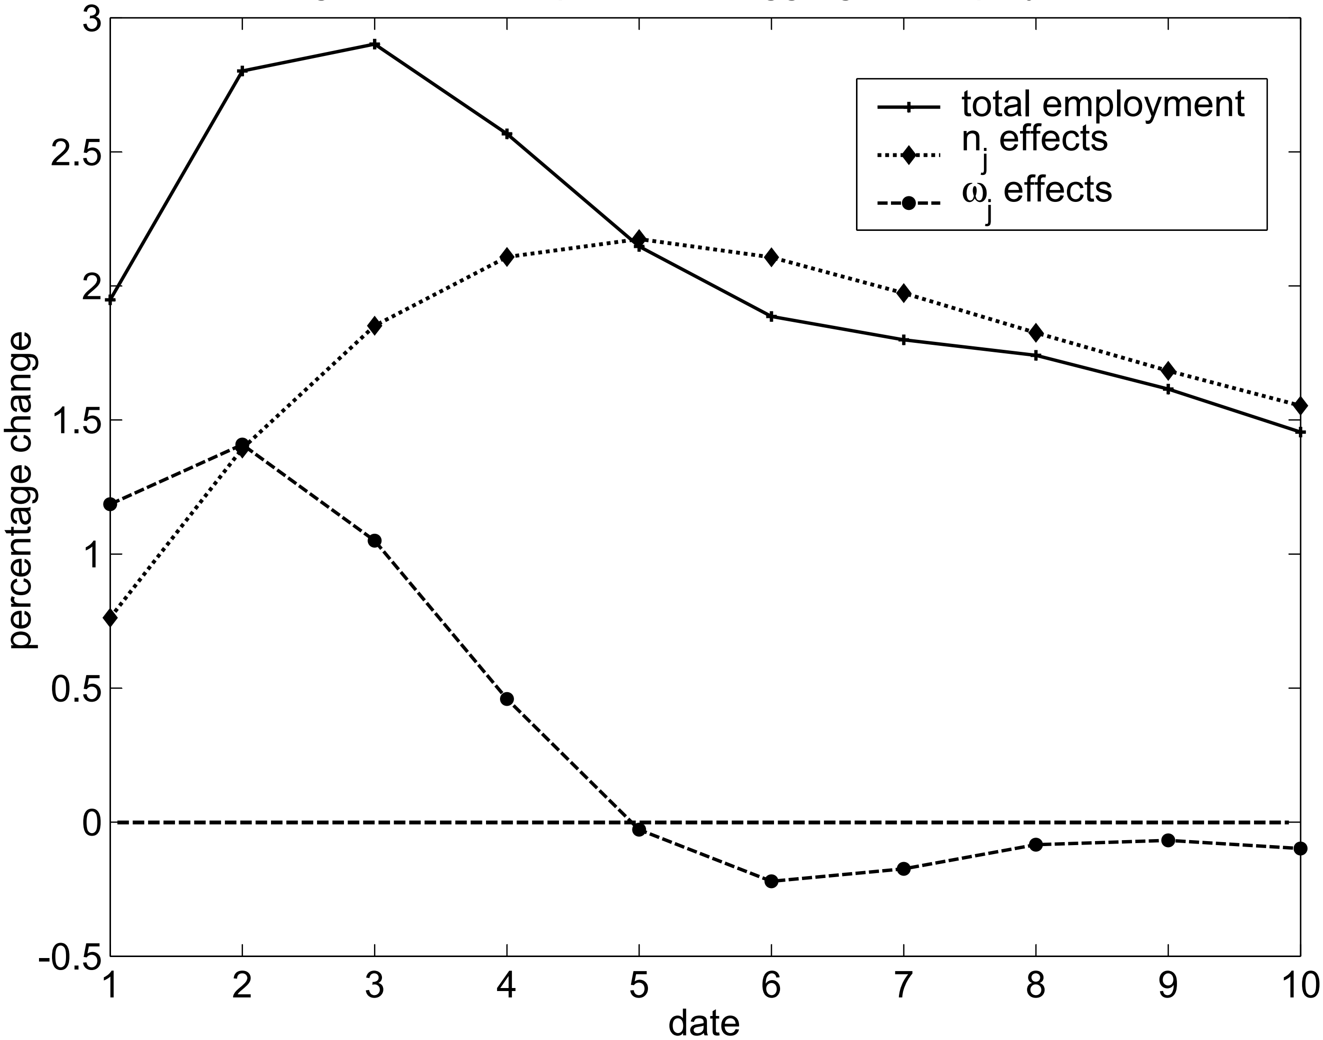

Figure 5: Effects of Equilibrium