

Federal Reserve Bank of Minneapolis
Research Department Staff Report 327

Revised August 2004

Partial Adjustment Without Apology*

Robert G. King

Boston University, Federal Reserve Bank of Richmond,
and National Bureau of Economic Research

Julia K. Thomas

University of Minnesota, Federal Reserve Bank of Minneapolis,
and National Bureau of Economic Research

ABSTRACT

Many kinds of economic behavior involve 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 such microeconomic adjustment is well described by a new form of partial adjustment model that aggregates the actions of heterogeneous producers.

We develop a model where individual establishments infrequently alter the sizes of their workforces because such adjustments involve fixed costs. In the market equilibrium, employment responses to aggregate disturbances include changes both in target employments 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 discrete adjustment models, our *generalized partial adjustment model* is sufficiently tractable to allow general equilibrium analysis, and it naturally extends to accommodate persistent differences in productivity across establishments in general equilibrium.

* Please direct correspondence to Julia Thomas, Department of Economics, University of Minnesota, 271 19th Avenue South, Minneapolis, MN 55455; jkt@econ.umn.edu. We thank Randall Wright and two anonymous referees for their comments. Any remaining errors are our own. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or Minneapolis, or the Federal Reserve System.

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 that are increasing in the 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

We present a *generalized partial adjustment* model in which individual production units adjust in a discrete and occasional manner, yet there is smooth adjustment at the aggregate level. Specifically, individual units face differing fixed costs of adjustment, so the timing of their adjustments is infrequent and asynchronized while aggregation across plants leads to a smooth pattern of aggregate factor demand well-approximated by the standard partial adjustment model. Our exposition of this model's relation to the traditional model commonly used in empirical work is unique to this paper.

Our basic framework is sufficiently tractable that it has already been applied to ex-

¹See, for example, Kollintzas (1985).

amine several topics, among them price adjustment and capital investment.² Here, we apply it to employment which, relative to the investment application, requires a different timing to trace the resulting distribution of production. We provide the first comprehensive presentation of the framework so that researchers may conveniently adapt it to study other problems. To facilitate its broad application, we then extend the method to allow for persistent idiosyncratic shocks.

Our model provides 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, with adjustment probabilities being 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.³

One key stylized fact uncovered in the empirical literature is that an important route through which aggregate shocks affect aggregate employment is by changing the fraction of plants that choose to adjust. Accordingly, we develop a model where the aggregate adjustment rate is an endogenous function of the state of the economy. While our generalized model is not observationally equivalent to the traditional partial adjustment model with time-invariant aggregate adjustment rates, impulse responses establish that it retains

²Dotsey, King and Wolman (1999) use the framework to analyze price adjustment; Thomas (2002) uses it to examine investment.

³Generalized (S, s) models were first studied by Caballero and Engel (1999) to explain the observed lumpiness of plant-level investment demand.

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. Moreover, it is sufficiently tractable to accommodate additional sources of heterogeneity; thus, beyond achieving consistency with the stylized facts highlighted here, the framework has the potential for use in applications designed to examine richer aspects of establishment-level dynamics.

The organization of this discussion is as follows. Section 2 briefly reviews the essential properties of the standard partial adjustment model, and section 3 describes the evidence on microeconomic adjustment patterns that the standard model fails to explain. Section 4 develops a model that is consistent with the observation that individual establishments hire varying amounts of labor at discrete and occasional times, and it illustrates a resulting hedging effect on the demand for labor. Next, section 5 embeds the framework within a fully specified general equilibrium macroeconomic model and endogenizes 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 our *generalized partial adjustment* model using a series of numerical examples.⁴ A distinguishing feature of the model, beyond its consistency with the evidence on employment adjustment in section

⁴Our 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 our model 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. Our model can also generate such rich dynamics, although it does not do so under the parameters selected here.

3, is that it has the potential to reproduce the sharp changes in market employment demand found in the data during episodes involving large changes in productivity.⁵ Moving from our market demand examples, we then provide counterpart results that illustrate the role of equilibrium in shaping the aggregate response to shocks. Finally, section 6 illustrates how the framework is tractably extended to allow for persistent differences in productivity across establishments, and section 7 concludes.

2 The standard partial adjustment model

The standard partial adjustment model relates current employment, N_t , to *target* or desired employment, N_t^* , through $N_t - N_{t-1} = \kappa[N_t^* - N_{t-1}]$, where $\kappa \in (0, 1)$ is the fraction of the gap closed in the period. This specification implies the influence of past actual or desired employment on current employment,

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

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. In the absence of costly adjustment, assume that the firm's workforce declines at the rate $d \in [0, 1)$ due to quits or mismatches. If e_t employees are hired at time t , then $N_t = (1 - d)N_{t-1} + e_t$, and the cost of the workforce adjustment is $\Xi(e_t) = \frac{B}{2}e_t^2$, where $B > 0$.⁶ Let z_t reflect current productivity, and w_t be the real wage (both serially correlated random variables known at date t), and let production be

⁵This 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.

⁶This captures the idea that the firm's marginal adjustment cost is rising in the extent of employment adjustment; the same idea is incorporated in alternative adjustment cost functions used in applied work.

$f(N_t, z_t)$. Discounting future earnings by $\beta \in (0, 1)$, the firm chooses $\{N_t, e_t\}_{t=0}^{\infty}$ to maximize its expected present discounted value, $\mathbf{E}\left[\sum_{t=0}^{\infty} \beta^t \left(f(N_t, z_t) - \Xi(e_t) - w_t N_t\right) \mid z_0, w_0\right]$, subject to $N_t = (1 - d)N_{t-1} + e_t$ and given initial employment N_{-1} . If the production function is quadratic in employment, it is straightforward to show that

$$N_t^* = \left[E_t \sum_{j=0}^{\infty} (\beta/\kappa)^j (\chi_a a_{t+j} - \chi_w w_{t+j})\right], \quad (2)$$

demonstrating that the presence of lags in employment implies leads under rational expectations, as stressed by Sargent (1978).⁷

The key implications of the model are: (i) current employment, N_t , is directly related to lagged employment, N_{t-1} , because adjustments are costly, and (ii) expectations of future wages and productivity influence current employment through the target, N_t^* , because, given adjustment costs, its choice will in part determine future employment. Taken together, these features imply that adjustment costs dampen the response to changes in current wage and productivity and yield smooth, gradual changes in employment over time.

3 Disconcerting evidence

While the traditional partial adjustment model offers a tractable framework with 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 suggests a number of stylized facts about individual and aggregate adjustment that we summarize here.

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

⁷In (1) and (2), the parameters κ , χ_a and χ_w depend on the adjustment cost parameter B , the discount factor β and the parameters of the production function.

over the sample. Employment exhibits long periods of constancy broken by infrequent and 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 resemble the dynamics of aggregate output and appear consistent with smooth adjustment behavior of aggregates. More specifically, he 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.⁸

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 in a particular date is typically called an *adjustment hazard*. Caballero and Engel (1993) construct a general framework for studying aggregate employment changes that can incorporate a variety of assumptions about how adjustment hazards are related to aggregate conditions. Using U.S. manufacturing data from 1961 through 1983, Caballero and Engel examine the dynamics of aggregate employment changes under

⁸Hamermesh compares log likelihood values from the estimation of a smooth adjustment model based on quadratic costs to those from a lumpy adjustment fixed-cost alternative. For plant level data, the latter achieves much larger likelihood values, indicating that lumpy adjustment based on fixed costs better describes the micro data. Further, his switching model estimates of the percentage ‘disequilibrium’ required to induce adjustment are large, suggesting that plants vary employment with a non-marginal adjustment only in the presence of substantial shocks to expected output. However, differences at the aggregate level are too small to discriminate between models, as is the case when they are compared using 4-digit SIC data. Thus, lumpy adjustment behavior at the microeconomic level is obscured by aggregation. From this and similar evidence, Hamermesh and Pfann (1996, page 1274) 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.”

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, reflecting 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 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.

4 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 stylized facts discussed above. The framework we develop exemplifies such a model. In particular, it delivers the implication

that, while individual establishments' employment adjustments are discrete, (fact 1), their asynchronous timing implies a smooth aggregate employment series similar to that implied by the traditional partial adjustment model, (fact 2). Moreover, 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), and the model can produce substantial responses of employment to aggregate shocks without relying on any shifts in the distribution of adjustment costs, (fact 5). At the same time, it can be readily incorporated into a general equilibrium model, so that the relationship between adjustment hazards and macroeconomic conditions can be studied.

We assume a large and fixed number of units, each making discrete choices about their employment adjustment over time. Production at the plant level is constant returns in labor and a fixed input, which we normalize to 1, $f(n_t, 1, z_t)$.⁹ Any unit that does not adjust its workforce sees it decay at rate d ,

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

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 to allow consistency with stylized fact 3 in section 5.

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 adjustment 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 $\alpha = [\alpha_1, \alpha_2, \dots, \alpha_{J-1}, 1]$ over time. We further assume that $\alpha_{j-1} < \alpha_j$ for all $j = 1, 2, \dots, J$,

⁹The presence of the fixed input allows determination of the employment choice at the production unit.

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_{jt} represent the current labor stock of a production unit that last adjusted its employment j periods ago.¹⁰ We use the notation $V_j(n_{jt}, z_t, w_t)$ to denote the value of a production unit that last adjusted j periods ago, entering the current period with a workforce of n_{jt} , that is not currently able to adjust its employment, and use $V_0(z_t, w_t)$ to denote the value of a production unit currently able to adjust. For a unit that is currently readjusting its stock of labor,

$$V_0(z_t, w_t) = \max_{n_{0t}} \left(f(n_{0t}, z_t) - w_t n_{0t} + \beta \mathbf{E} \left[\alpha_1 V_0(z_{t+1}, w_{t+1}) + (1 - \alpha_1) V_1((1 - d) n_{0t}, z_{t+1}, w_{t+1}) \mid z_t, w_t \right] \right), \quad (4)$$

where n_{0t} is freely chosen, and (z_t, w_t) follows a joint Markov process.¹¹ 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 allowing adjustments on other margins, such as in hours-per-worker. Their labor evolves according to $n_{jt} = (1 - d)n_{j-1,t-1}$, and their value functions obey the functional equation

$$V_j(n_{jt}, z_t, w_t) = f(n_{jt}, z_t) - w_t n_{jt} + \beta \mathbf{E} \left[\alpha_{j+1} V_0(z_{t+1}, w_{t+1}) + (1 - \alpha_{j+1}) V_{j+1}((1 - d) n_{jt}, z_{t+1}, w_{t+1}) \mid z_t, w_t \right]. \quad (5)$$

Adjusting production units choose employment so as to maximize the right-hand side of (4), which results in an efficiency condition of the following form:

¹⁰Except where necessary for clarity, we suppress commas in subscripts throughout this text.

¹¹As n_{0t} is our model's counterpart to target employment in the traditional model, we occasionally refer to it as n^* when making comparisons below.

$$D_1 f(n_{0t}, z_t) - w_t + \beta \mathbf{E} \left[(1 - \alpha_1) (1 - d) D_1 V_1 \left((1 - d) n_{0t}, z_{t+1}, w_{t+1} \right) \middle| z_t, w_t \right] = 0.$$

A notable feature of this condition is 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_{jt} 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 (5) above, we can determine the marginal value of additional workers:

$$\begin{aligned} D_1 V_j(n_{jt}, z_t, w_t) &= D_1 f(n_{jt}, z_t) - w_t \\ &\quad + \beta \mathbf{E} \left[(1 - \alpha_{j+1}) (1 - d) D_1 V_{j+1} \left((1 - d) n_{jt}, z_{t+1}, w_{t+1} \right) \middle| z_t, w_t \right]. \end{aligned}$$

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

$$D_1 f(n_{0t}, z_t) - w_t + \mathbf{E} \sum_{j=1}^{J-1} \left[[\beta(1-d)]^j \varphi_j \left(D_1 f((1-d)^j n_{0t}, z_{t+j}) - w_{t+j} \right) \middle| z_t, w_t \right] = 0, \quad (6)$$

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

$$\varphi_j \equiv \prod_{k=1}^j (1 - \alpha_k), \quad j = 1, \dots, J-1. \quad (7)$$

In the special case of a Cobb-Douglas production function, $y = zn^\gamma$, the condition in (6) may be explicitly solved for adjusting units' optimal labor demand as

$$n_{0t} = \left[\frac{E \sum_{j=0}^{J-1} \left[\gamma \beta^j (1-d)^{\gamma j} \varphi_j z_{t+j} \middle| z_t, w_t \right]}{E \sum_{j=0}^{J-1} \left[\beta^j (1-d)^j \varphi_j w_{t+j} \middle| z_t, w_t \right]} \right]^{\frac{1}{1-\gamma}},$$

which depends positively on current and expected future productivity and negatively on current and expected future wages. Since this forward-looking labor demand is similar to the behavior of target employment in the standard partial adjustment model of section 2, we sometimes refer to it in this manner. However, it is worth stressing that the economic reasons for this are somewhat different. In the standard model, the firm's labor demand is forward-looking because current adjustments affect the future costs that a firm encounters when it adjusts. Here, by contrast, labor demand is forward-looking because a current adjuster is aware that it may not be able to adjust its employment again in the near future.

The condition in (6) also implies that our generalized adjustment model has a hedging property arising because of forecasted future labor force departures. In particular, if wages and productivities are expected to be constant over time, then an establishment will demand more employment than it would in a frictionless environment. When changes in wages and productivities are suppressed, the target employment level solving (6) is a constant $n^*(\alpha, z, w)$. Let n^s represent the static optimum satisfying $D_1 f(n^s, z) - w = 0$ that would be chosen if the unit could adjust its employment in every period with certainty. Given concavity of f , it follows that $[D_1 f((1-d)^j n_0, z) - w] < [D_1 f((1-d)^{j+1} n_0, z) - w]$. This implies the summation in (6) evaluated at $n_0 = n^s$ is strictly positive. Moreover, as both this sum and its preceding expression, $D_1 f(n_0, z) - w$, are decreasing in n_0 , the dynamic optimum, n^* , must exceed the static optimum, n^s .

Production units hire more labor than they currently need in an effort to hedge against the possibility that they may be unable to hire in the immediate future. Further, n^* will be larger the higher is this probability of future nonadjustment; for instance, given d and $\alpha_2 \cdots \alpha_{J-1}$, a reduced probability of adjustment in the first period after an adjustment, (lower α_1) yields higher values for $\varphi_1, \dots, \varphi_{J-1}$ and thus a higher value for the summation at any n_0 . The higher is the probability of being unable to restock employment, the stronger is the hedging motive.

5 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.¹² Within each date, any individual production unit faces a random cost ξ that it must pay in order to adjust its employment. This cost is drawn from a time-invariant distribution over $[0, B]$ that is summarized by the CDF $G(\xi)$ and associated PDF $g(\xi)$.

In addition to endogenizing adjustment timing, we also illustrate that it is straightforward to incorporate the generalized partial adjustment approach into a general equilibrium setting, which we do now by imposing two restrictions on the prices faced by establishments within each date. First, asset-market clearing will require that all establishments discount their future profit flows by households' marginal rate of substitution between current and future consumption denoted here by $\beta \frac{p_{t+1}}{p_t}$. Equivalently, establishments value their current output by p_t , the current marginal utility of consumption, and discount their future values by the household subjective discount factor β . Next, the equilibrium wage, w_t , will equal households' marginal rate of substitution between current leisure and consumption, $\frac{D_2 u(c, 1-N)}{D_1 u(c, 1-N)}$. Provided that these restrictions are satisfied, the role of households in the economy is effectively subsumed, and equilibrium allocations are retrieved as the aggregate of establishments' decisions.¹³

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

¹²The generalized adjustment model developed here 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 (2003a) investigate investment dynamics and Khan and Thomas (2003b) use a similar approach to study (S,s) inventory accumulation. In this study, we use linear approximation methods in the tradition of Sargent (1978) to explore the general equilibrium dynamics, as do Dotsey, King and Wolman (1999) and Thomas (2002).

¹³See Khan and Thomas (2003a) for further explanation.

elapsed since the last active employment adjustment. Given its current cost draw ξ , and given its start of period employment, n_{jt} , and S_t , the aggregate state of the economy determining prices and expectations, such an establishment will adjust its employment if its fixed cost does not exceed the value of the adjustment, that is, if $V_0(S_t) - V_j(n_{jt}, S_t) \geq \xi$.¹⁴ 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 $\bar{\xi}_{jt}$ such that

$$\bar{\xi}_{jt} = V_0(S_t) - V_j(n_{jt}, S_t). \quad (8)$$

All production units in the j^{th} time-since-adjustment group with adjustment costs at or below the threshold in (8) will choose to adjust. As a result, the fraction of plants adjusting out of any particular group j , $j = 1, \dots, J - 1$, is given by

$$\alpha_{jt} = G(\bar{\xi}_{jt}). \quad (9)$$

From (8), note that these adjustment fractions are functions of the plant-level state vector, (n_{jt}, S_t) . We assume that the stochastic processes for productivity, wages and interest rates are such that, given the function f and the discount factor β , $B < V_0(S_t) - V_J(n_{Jt}, S_t)$ for all values of the vector (n_{Jt}, S_t) . This assumption, which follows naturally from $B < \infty$, given bounded processes z_t , w_t and p_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 (8) - (9). With state-dependent probability $\alpha_{j+1,t+1}$,

¹⁴As will be made explicit below, S_t includes two endogenous vectors that together identify the start-of-period distribution of plants over employments, alongside exogenous aggregate productivity, z_t . We assume z follows a Markov process that is taken as given by all agents, as is the evolution of the endogenous aggregate state, A , according to a mapping $A' = \Psi(A, z)$ that, in equilibrium, results from the aggregation of individual actions.

a production unit entering period $t + 1$ in group $j + 1$ will adjust at that date. The counterpart to (4), the value of a plant that is currently adjusting its labor, is

$$V_0(S_t) = \max_{n_{0t}} \left(f(n_{0t}, z_t) - w_t n_{0t} + \beta \mathbf{E} \left[\frac{p_{t+1}}{p_t} \left[\alpha_{1,t+1} V_0(S_{t+1}) - \xi_{1,t+1} + (1 - \alpha_{1,t+1}) V_1((1-d)n_{0t}, S_{t+1}) \right] \mid S_t \right] \right), \quad (10)$$

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, $\xi_{1,t+1} = \int_0^{G^{-1}(\alpha_{1,t+1})} xg(dx)$. Similarly, the value of a plant that last adjusted j periods ago, the counterpart to equation (5), is

$$V_j(n_{jt}, S_t) = f(n_{jt}, z_t) - w_t n_{jt} + \beta \mathbf{E} \left[\frac{p_{t+1}}{p_t} \left[\alpha_{j+1,t+1} V_0(S_{t+1}) - \xi_{j+1,t+1} + (1 - \alpha_{j+1,t+1}) V_{j+1}((1-d)n_{jt}, S_{t+1}) \right] \mid S_t \right]. \quad (11)$$

Adjusting plants exit the j^{th} group for the adjustment group and choose an optimal employment level n_{0t} (n_t^*) satisfying the marginal profit condition below, which generalizes (6) to reflect optimal adjustment probabilities:

$$D_1 f(n_{0t}, z_t) - w_t + \mathbf{E} \sum_{j=1}^{J-1} \left[\frac{p_{t+j}}{p_t} [\beta(1-d)]^j \varphi_{j,t+j} [D_1 f((1-d)^j n_{0t}, z_{t+j}) - w_{t+j}] \mid S_t \right] = 0. \quad (12)$$

Here, as in (7), $\varphi_{j,t+j}$ is the probability the unit will make no further adjustment in the next j periods. That is, for $j = 1, \dots, J - 1$,

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

In practice, it is convenient to break the large forward-looking condition determining target employment into J first-order stochastic difference equations as follow:

$$D_1 f(n_{0t}, z_t) - w_t + \beta(1 - d)E\left[\frac{p_{t+1}}{p_t}\Omega_{1,t+1} \mid S_t\right] = 0, \quad (13)$$

where, for $j = 1, \dots, J - 1$,

$$\Omega_{jt} \equiv (1 - \alpha_{jt})\left(D_1 f(n_{jt}, z_t) - w_t + \beta(1 - d)E\left[\frac{p_{t+1}}{p_t}\Omega_{j+1,t+1} \mid S_t\right]\right). \quad (14)$$

5.1 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 $\theta_t = [\theta_{1t}, \dots, \theta_{Jt}]$, with each θ_{jt} representing the fraction of units that begin the period having last adjusted j periods prior to the current date.¹⁵ Let $\omega_{0t} \equiv \sum_{j=1}^J \theta_{jt}\alpha_{jt}$ denote total adjusting units in any date t .¹⁶ The elements of this vector are then as follow:

$$\theta_{1t} = \omega_{0,t-1} \quad (15)$$

¹⁵More 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 $[n_{t-1}^*, \dots, n_{t-J}^*]$ from which the current support is trivially retrieved. Note that the time-since-adjustment approach to tracking the plant distribution pursued here allows us to capture the time-varying distribution of establishments over employment levels using a linear systems solution approach. We could instead directly track the measure associated with each possible employment level. However, in that case, we would need to include employments that at times have zero population, necessitating a nonlinear solution method as, for instance, in the investment study of Khan and Thomas (2003a).

¹⁶Given a fixed measure of production units, this overall adjustment rate is $\omega_{0t} = 1 - \sum_{j=1}^{J-1} \omega_{0,t-j}\varphi_{jt}$.

$$\theta_{jt} = (1 - \alpha_{j-1,t-1})\theta_{j-1,t-1} \quad \text{for } j = 2, \dots, J. \quad (16)$$

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:

$$N_t = n_t^* \sum_{j=1}^J \theta_{jt} \alpha_{jt} + \sum_{j=1}^{J-1} \theta_{jt} (1 - \alpha_{jt}) d^j n_{t-j}^*. \quad (17)$$

Here, n_t^* is the target value of employment that solves (12); in the case of Cobb-Douglas

production, this is $n_t^* = \left[\frac{E \sum_{j=0}^{J-1} \left[\gamma \beta^j (1-d)^{\gamma j} \varphi_{j,t+j} z_{t+j} | z_t, w_t \right]}{E \sum_{j=0}^{J-1} \left[\beta^j (1-d)^j \varphi_{j,t+j} w_{t+j} | z_t, w_t \right]} \right]^{\frac{1}{1-\gamma}}$.

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 employments, as in the traditional partial adjustment model (1). Consequently, while the underlying production unit level demands are adjusted discretely and occasionally, the market demands vary smoothly in every time period. Further, since each target employment, n_{t-j}^* , $j = 1, \dots, J - 1$, involves expectations of future wages and productivities, so does market labor demand.

While equation (17) 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 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 model, our economywide rate of adjustment responds to changes in aggregate conditions, including changes in economic policy.

5.2 Planning representation

The generalized partial adjustment model described above may be derived as the solution to a single dynamic optimization problem, which makes the link to the standard model of section 2 more direct. We briefly outline this reformulation to illustrate the tractability of the approach and thus its suitability for applications.¹⁷ While we rely on the equivalence between a social planning and competitive equilibrium solution in this section, as in Lucas and Prescott (1971), it is important to stress that the generalized partial adjustment approach can also be applied to settings in which competitive equilibrium is not optimal.¹⁸

The aggregate representation consolidates the ownership of all plants, differentiated by their time since last adjustment, $j = 1, \dots, J$, into a single entity, a planner acting to maximize the expected discounted lifetime utility of a representative household. Using the notation $\boldsymbol{\theta}_t \equiv [\theta_{1t}, \dots, \theta_{Jt}]$, $\mathbf{n}_t \equiv [n_{1t}, \dots, n_{Jt}]$, and $\boldsymbol{\alpha}_t \equiv [\alpha_{1t}, \dots, \alpha_{Jt}]$ to describe the economywide distribution of plants, employment, and adjustment fractions across groups, the planner's total available output is

$$Y_t = f(n_{0t}, z_t) \sum_{j=1}^J \theta_{jt} \alpha_{jt} + \sum_{j=1}^{J-1} \theta_{jt} (1 - \alpha_{jt}) f(n_{jt}, z_t). \quad (18)$$

Total employment is an analogous sum of the employments of adjusting and non-adjusting establishments,

$$N_t^D = n_{0t} \sum_{j=1}^J \theta_{jt} \alpha_{jt} + \sum_{j=1}^{J-1} \theta_{jt} (1 - \alpha_{jt}) n_{jt}. \quad (19)$$

Finally, economywide adjustment costs are

¹⁷Here, 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 in application.

¹⁸In its application to the analysis of price adjustment by Dotsey, King and Wolman (1999), for example, the presence of monopolistic competition means that equilibrium is not optimal.

$$Q_t = \sum_{j=1}^J \theta_{jt} \Gamma(\alpha_{jt}), \quad (20)$$

where $\Gamma(\alpha) = \int_0^{G^{-1}(\alpha)} xg(dx)$ is the total volume of costs averaged across plants in a group if fraction α of that group adjusts.

Given the current distribution of plants over time-since-last-adjustment groups, the associated employment levels, and aggregate productivity, the planner chooses fractions of plants adjusting $(\alpha_{jt})_{j=1}^{J-1}$ and optimal employment for those that are adjusting their workers, n_{0t} , which together determine the next period distribution of plants, $\boldsymbol{\theta}_{t+1}$, and the household's current consumption and work hours. The planner's problem is

$$\begin{aligned} W(\boldsymbol{\theta}_t, \mathbf{n}_t, z_t) = & \max_{\Lambda_t} u(c_t, 1 - N_t) + \beta EW(\boldsymbol{\theta}_{t+1}, \mathbf{n}_{t+1}, z_{t+1} \mid \boldsymbol{\theta}_t, \mathbf{n}_t, z_t) \quad (21) \\ & + \lambda_t [Y_t - Q_t - c_t] \\ & + w_t \lambda_t [N_t - N_t^D], \\ & + s_{0t} \lambda_t \left[\sum_{j=1}^J \theta_{jt} \alpha_{jt} - \theta_{1,t+1} \right] \\ & + \sum_{j=1}^{J-1} s_{jt} \lambda_t \left[\theta_{jt} (1 - \alpha_{jt}) - \theta_{j+1,t+1} \right] \end{aligned}$$

subject to $n_{j+1,t+1} = (1 - d)n_{jt}$, for $j = 0, \dots, J - 1$, and subject to (18)-(20), where $\Lambda_t = [c_t, N_t, n_{0t}, \{\alpha_{jt}\}_{j=1}^{J-1}, \{\theta_{j+1,t+1}\}_{j=0}^{J-1}]$.

The solution to this problem will satisfy the constraints above with equality and a series of efficiency conditions that follow. First, the standard conditions apply to the choice of household consumption and labor supply,

$$\begin{aligned} \lambda_t &= D_1 u(c_t, 1 - N_t) \\ w_t \lambda_t &= D_2 u(c_t, 1 - N_t). \end{aligned}$$

From these two equations, it is clear that the output price, p_t , and the real wage, w_t , faced by establishments in the decentralized economy examined above must correspond to the

multipliers λ_t and w_t , respectively, if the competitive allocation is to match that obtained here.

Note that the multipliers s_{jt} attached to the distributional constraints in (21) represent date t post-production valuations of establishments that will enter the next date in plant group $j + 1$. To clarify the equivalence between the planning allocation and that in the decentralized economy, we define the pre-production valuations of establishments as

$$v_{jt} \equiv f(n_{jt}, z_t) - w_t n_{jt} + s_{jt}, \text{ for } j = 0, \dots, J - 1,$$

and we use these, rather than the original multipliers, in representing the optimal adjustment fractions. Efficiency with respect to the choice of α_{jt} requires that the solution to this problem satisfy

$$G^{-1}(\alpha_{jt}) = v_{0t} - v_{jt},$$

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 α_{jt} , $j = 1, \dots, J - 1$, and is equivalent to (8) provided the multipliers v_{jt} attain the same value as before. That this is the case may be seen from the efficiency conditions with respect to $\theta_{j+1,t+1}$, $j = 0, \dots, J - 1$, which imply that the value associated with a plant with employment level n_{jt} satisfies

$$v_{jt} = f(n_{jt}, z_t) - w_t n_{jt} + \beta \mathbf{E} \left[\frac{\lambda_{t+1}}{\lambda_t} \left(\alpha_{j+1,t+1} v_{0,t+1} - \Gamma(\alpha_{j+1,t+1}) \right. \right. \\ \left. \left. + (1 - \alpha_{j+1,t+1}) v_{j+1,t+1} \right) \mid \boldsymbol{\theta}_t, \mathbf{n}_t, z_t \right].$$

These expressions are equivalent to the plant Bellman equations of section 5, since the expected adjustment cost conditional on adjustment in (11) is equal to $\Gamma(\alpha_{j+1,t+1})$, the average cost paid by adjusting plants, by definition of $\Gamma(\cdot)$. Finally, the efficiency condition with respect to the choice of n_{0t} may be expressed as (13)-(14), provided $p_t = \lambda_t$ at every date. Therefore, the solution to the planning problem, given the aggregate state $(\boldsymbol{\theta}_t, \mathbf{n}_t, z_t)$, is the same as in the decentralized economy of the previous section.

5.3 Numerical examples: The five stylized facts

We use a series of numerical examples to illustrate several interesting properties of the model developed above, and to contrast its dynamics to those of the traditional 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 examples involve functional forms and parameter values that are standard; production at the plant level is described by a Cobb-Douglas production function $f(n, z) = zn^\nu$ with $\nu = 0.66$. 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 plant'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).

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 traditional model, we assume the quadratic cost parameter is $B = 4$. 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, 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 plants across groups is necessarily downward sloping, given the law of motion for θ in (16).

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 from their steady state values, in response to a persistent rise in aggregate productivity, for the 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 model with a fixed vector of time-dependent adjustment fractions. Finally, SD denotes the response in the generalized state-dependent partial adjustment model. There, fixed costs of adjustment dissuade some production units from responding immediately 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 traditional partial adjustment model. This hump-shaped response in employment, most pronounced for the SD model, is absent in a frictionless model of employment adjustment. There, without adjustment costs, the shape of the employment response is identical to the monotonic response of the auto-correlated productivity shock.

The TD 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, 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 its current adjustors.

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, allowing for sharper changes in market employment, 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 is present when adjustment rates are held fixed, it is apparent that they arise due to 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 percentage changes in extensive margin adjustment through changes in the fraction of production units adjusting, $\omega_{0t} = \sum_{j=1}^J \theta_{jt} \alpha_{jt}$, 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 stylized fact 3 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, in this example, changes in the number of adjusting plants are more important than changes in the employment chosen by such plants in determining movements in aggregate employment. Furthermore, the latter is responsible for the cyclical pattern seen in figure 2 for the aggregate series.

Comparing panels (a) and (b) of figure 3, note that, while the target employment response monotonically declines, 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 large shift in the distribution of production units away from higher time-since-adjustment groups and into group 1 starting the next period. Given rising adjustment hazards, only a small fraction of these extra members find it worthwhile to adjust again, 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 resettles.

Figure 4 aggregates the effects of changes in intensive margin versus extensive margin adjustment to provide a decomposition of the market employment response into two underlying components: “ n_j effects” associated with changes in employment levels across

groups (due to changes in target employments) and “ ω_j effects” arising from changes in the distribution of plants across these groups at the time of production, $\omega_{jt} \equiv (1 - \alpha_{jt})\theta_{jt}$, $j = 1, \dots, J$, (due to changes in the fractions adjusting from each group). Specifically, at each date, the percentage deviation from steady state in aggregate employment is given by

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

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 \hat{n}_{jt} and $\hat{\omega}_{jt}$ represent percentage 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 , $j = 1, \dots, J$. 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.

5.4 General equilibrium effects

One of the key features of our approach is that discrete micro-level adjustment dynamics can be readily introduced into a general equilibrium setting. So far, we have used the dynamic model to study the influence of variations in productivity on aggregate labor demand and the adjustment decisions of individual units, holding fixed the real wage rate and the real interest rate. A general equilibrium model allows changes in productivity to affect the wage and interest rates, so that its dynamics are more complicated. For example, when a rise in productivity increases labor demand, this will bring about some increase in the wage rate in order to clear the labor market. This wage change has implications for both the level of labor that adjusting establishments select (the target solving (12)) and the fractions of establishments that choose to adjust from each current employment (the adjustment fractions solving (8)-(9)). Our equilibrium analysis is designed to be very simple, but it illustrates some important points. Here, we assume a particular functional form for the representative household's preferences, so as to generate restrictions on the behavior of the wage rate and the interest rate.¹⁹

Figure 5 compares the dynamic general equilibrium response to an aggregate productivity shock within the state-dependent generalized adjustment model to the response arising without market-clearing variations in wages and interest rates. Quantitatively, as might be expected, equilibrium price movements sharply dampen the response in employment, and

¹⁹In particular, maintaining the functional forms and parameter values assumed above, we assume that the representative household's momentary utility function is $U(C, N) = \log\left(C - \chi \frac{N^{1+\gamma}}{1+\gamma}\right)$, where $\chi = 2.55$ and $\gamma = 0.50$. This specification implies that there is a steady-state level of labor of $n = 0.20$ and a general equilibrium labor supply elasticity of $\gamma^{-1} = 2$ with respect to the real wage rate. Higher values of γ would imply sharper differences between the equilibrium and fixed price models, as these would raise the responsiveness of the wage to changes in employment demand.

hence output, to a persistent change in productivity. However, in contrast to the investment analysis of Thomas (2002), equilibrium does not eliminate the influence of costly and discrete adjustment. In particular, the level of employment continues to display a hump in the dynamic response due to these costs, which is also an implication of the standard partial adjustment model discussed in section 2 above.

There are also qualitative changes in both the extensive and intensive margins of employment adjustment with equilibrium movements in wages and interest rates. First, the previous nonmonotonicity in the fraction of units adjusting essentially disappears. This is because equilibrium price changes offset much of the large rise in target employment that would otherwise occur at the impact of the shock. With the rise in target employment dampened, establishments have less incentive to pay fixed costs to move up the timing of their employment adjustments. Thus, equilibrium reduces the jump in the total fraction adjusting, thereby reducing the disruptions to current and future plant distributions that cause these oscillations. Second, the smooth mean reversion in target employment becomes less regular. Nonetheless, target employment continues to be monotonic, and it is the changes in adjustment timing that lead to the hump-shaped aspects of the aggregate quantity responses. More establishments do choose to adjust employment when there is a favorable productivity shock, but they do not all adjust immediately.

6 Persistent idiosyncratic shocks

To this point, plants have been differentiated by only two features of their circumstances: (i) they have different realizations of adjustment costs, and hence (ii) they enter the period with different values of the endogenous state variable n_{jt} . However, there is ample evidence that establishments are affected by additional persistent plant-level states, such as stochastic variations in productivity. In this section, we show how our equilibrium generalized partial adjustment model is tractably extended to allow for persistent productivity shocks to establishments. The methods, however, can be applied to other persistent

exogenous micro states, such as variations in product demand for monopolistic competitors or shifts in the distribution of adjustment costs.

We assume that the plant-specific productivity shocks follow an M -state Markov process; $a \in \{a_1, \dots, a_M\}$ with transition probabilities given by the time-invariant matrix Φ ; specifically, the probability of transiting from state a_l to state a_m is given by $\phi(l, m)$, for $l = 1, \dots, M$ and $m = 1, \dots, M$. We begin by defining some notation that will describe the distribution of establishments over employment and productivity at each date, then show how the aggregation is handled, and next proceed to outline the associated planning problem. For brevity, we omit the corresponding decentralized representation of the economy, although the mapping should be transparent by comparison to the sections above.

At the start of date t , any establishment is identified by its current productivity draw, a , and its current employment level. We continue to assume that, when not actively adjusted, a plant's employment declines at rate d across dates, and active adjustments to its stock incur a fixed cost, ξ , drawn from the time-invariant distribution $G(\xi)$ with associated PDF $g(\xi)$. Given the effect of current plant-specific productivity draws on adjustment decisions and on the target employments selected by current adjustors (and hence on future distributions), the economy's aggregate state, S , will now include $M^2 + M$ time-since-adjustment vectors that together describe the current start-of-period distribution of establishments over labor and productivity. We preserve the ability to solve this economy linearly by tracking the distribution of plants according to their membership in groups identified by (i) time-since-last adjustment, (ii) productivity draw at the date of last adjustment, and (iii) current productivity draw. As before, employment selected by adjusting establishments does not depend upon the current stock; however, it does depend upon current productivity.

To study the evolution of plant-level conditions, for each h and each l , we define $\theta_{jt}(h, l)$ as the start-of-date measure of plants that last adjusted j periods in the past to a target employment consistent with a_h (their productivity at the time of the adjustment) and that

have current productivity level a_l . Let $\alpha_{jt}(h, l)$ denote the corresponding fractions of each of these groups undertaking active employment adjustment within the current period. While adjustment fractions reach 1 within some finite number of periods, the full adjustment horizons for plants now depend upon the productivity they had when they last actively changed their employment and on their current productivity. Let $J(h, l)$ denote the full adjustment horizon associated with plants associated with productivity a_h at the time of last adjustment and current productivity a_l .²⁰ Because establishments transit across current productivities from date to date, each vector $\boldsymbol{\theta}_t(h, l) = [\theta_{jt}(h, l)]$ will have length $J_h \equiv \max\{J(h, 1), \dots, J(h, M)\}$. Finally, each of these vectors is associated with the vector of start-of-date employment levels $\mathbf{n}_t(h) = [n_{jt}(h)]$ of length J_h .²¹

6.1 Aggregation

The evolution of the plant distribution may be summarized as follows. First, there are M^2 equations representing the fractions of establishments that are current adjusters, and hence will begin the next period with time-since-last adjustment 1. One such equation holds for each current productivity, $l = 1, \dots, M$, and for each next-period productivity, $m = 1, \dots, M$. Each represents the fraction of all establishments that have current productivity a_l and adjust from their start-of-period employment level to the associated target $n_{0t}(l)$, and that will then enter next period identified by $(n_{1t}(l), a_m)$:

²⁰As with J in the model above, the horizons here, $J(h, l)$, are endogenous variables recovered in the solution for the economy's steady state. One cost of pursuing a linear systems solution for the dynamics is that we must assume that the economy stays sufficiently local to the steady state that these horizons are impervious to aggregate shocks. To know whether this assumption is reasonable in a given application, one must verify that all endogenous adjustment fractions remain strictly in the $(0, 1)$ interval at every date over long simulations.

²¹For example, in the case of a 2-state Markov shock, the start-of-date plant distribution over employment and productivities is completely summarized by four $\boldsymbol{\theta}$ vectors and two \mathbf{n} vectors. Using the notation outlined above, these are: $\boldsymbol{\theta}_t(1, 1)$ and $\boldsymbol{\theta}_t(1, 2)$, each of length $J_1 \equiv \max\{J(1, 1), J(1, 2)\}$; $\boldsymbol{\theta}_t(2, 1)$ and $\boldsymbol{\theta}_t(2, 2)$, each of length $J_2 \equiv \max\{J(2, 1), J(2, 2)\}$; $\mathbf{n}_t(1)$ of length J_1 ; $\mathbf{n}_t(2)$ of length J_2 .

$$\theta_{1,t+1}(l, m) = \phi(l, m) \sum_{h=1}^M \left(\sum_{j=1}^{J_h} \theta_{jt}(h, l) \alpha_{jt}(h, l) \right). \quad (22)$$

Next, there are M^2 sets of equations describing the non-adjusting population. Each set is identified by a particular (past, current) productivity combination, and each contains $J_h - 1$ separate equations, one for each possible time-since-last-adjustment. Specifically, each equation isolates the fraction of all plants that had productivity a_h , $h \in \{1, \dots, M\}$, at the time of their last adjustment j periods in the past, do not adjust this period (hence produce with employment $n_{jt}(h)$), and then draw productivity a_m , $m \in \{1, \dots, M\}$, at the start of the next period. These are the plants that will enter the next date having labor associated with productivity a_h from $j + 1$ periods in the past and having current productivity a_m :

$$\theta_{j+1,t+1}(h, m) = \sum_{l=1}^M \theta_{jt}(h, l) [1 - \alpha_{jt}(h, l)] \phi(l, m) \quad \text{for } j = 1, \dots, J_h - 1. \quad (23)$$

Finally, there are M sets of equations describing future employments of those that had productivity a_h , $h \in \{1, \dots, M\}$, at the time of their last adjustment:

$$n_{j+1,t+1}(h) = (1 - d)n_{jt}(h) \quad \text{for } j = 0, \dots, J_h - 1. \quad (24)$$

Equations (25)-(28) describe aggregate output gross of adjustment costs, aggregate labor demand and total adjustment costs in the economy. Aggregate output is total production by all establishments with current productivity $\{a_l\}_{l=1}^M$ that adjust to the optimal employment consistent with their productivity, together with the output of all nonadjustors (of each current productivity a_l) that last adjusted to an employment consistent with productivity $\{a_h\}_{h=1}^M$:

$$\begin{aligned} Y_t &= \sum_{l=1}^M \left[f(n_{0t}(l), a_l, z_t) \sum_{h=1}^M \sum_{j=1}^{J(h,l)} \theta_{jt}(h, l) \alpha_{jt}(h, l) \right] \\ &\quad + \sum_{l=1}^M \sum_{h=1}^M \left[\sum_{j=1}^{J(h,l)-1} f(n_{jt}(h), a_l, z_t) \theta_{jt}(h, l) [1 - \alpha_{jt}(h, l)] \right]. \end{aligned} \quad (25)$$

Total employment demand is an analogous sum of the employments of adjusting and non-adjusting establishments:

$$\begin{aligned}
N_t^D &= \sum_{l=1}^M \left[n_{0l}(l) \sum_{h=1}^M \sum_{j=1}^{J(h,l)} \theta_{jt}(h,l) \alpha_{jt}(h,l) \right] \\
&\quad + \sum_{l=1}^M \sum_{h=1}^M \left[\sum_{j=1}^{J(h,l)-1} n_{jt}(h) \theta_{jt}(h,l) [1 - \alpha_{jt}(h,l)] \right].
\end{aligned} \tag{26}$$

Finally, economywide adjustment costs are the total of those paid by establishments of each time-since-adjustment age that last adjusted to an employment consistent with productivity a_h and now have productivity a_l , again summing across past and current productivity levels:

$$Q_t = \sum_{l=1}^M \sum_{h=1}^M \left[\sum_{j=1}^{J(h,l)} \theta_{jt}(h,l) \Gamma(\alpha_{jt}(h,l)) \right], \tag{27}$$

where each $\Gamma(\alpha)$ in (27) represents the average adjustment cost paid per member in a given group, conditional on adjustment fraction α from that group,

$$\Gamma(\alpha) \equiv \int_0^{G^{-1}(\alpha)} xg(x). \tag{28}$$

6.2 Planning problem

The planning problem for the generalized partial adjustment model with persistent plant-specific productivities is listed below. Here, the aggregate state vector includes the $M^2 + M$ vectors that together describe the current distribution of establishments over employment and productivity, alongside current exogenous aggregate productivity, z :

$$S_t \equiv \left[[\boldsymbol{\theta}_t(h,l)]_{h,l=1}^M, [\mathbf{n}_t(h)]_{h=1}^M, z_t \right].$$

$$\begin{aligned}
W(S_t) &= \max_{\Lambda_t} u(c_t, 1 - N_t) + \beta EW(S_{t+1}|S_t) \\
&\quad + \lambda_t [Y_t - Q_t - c_t] + w_t \lambda_t [N_t - N_t^D],
\end{aligned} \tag{29}$$

subject to (22)-(28), where

$$\Lambda_t = \left\{ c_t, N_t, [n_{0t}(l)]_{l=1}^M, \left[[\alpha_{jt}(h, l)]_{j=1}^{J(h, l)-1} \right]_{h, l=1}^M, \left[[\theta_{j+1, t+1}(h, l)]_{j=0}^{J(h, l)-1} \right]_{h, l=1}^M \right\}.$$

The solution to this problem satisfies (22)-(28) and the constraints in (29) with equality, as well as a series of efficiency conditions that, after some algebra, may be written as follow.

First, aggregate consumption and labor supply satisfy

$$\begin{aligned} \lambda_t &= D_1 u(c_t, 1 - N_t) \\ w_t \lambda_t &= D_2 u(c_t, 1 - N_t). \end{aligned}$$

Next, we describe the conditions determining target employments. For each plant-specific productivity level a_l , define $\omega_{0t}(l)$ to be the total establishments that currently have this productivity and adjust their employment:

$$\omega_{0t}(l) \equiv \sum_{h=1}^M \sum_{j=1}^{J(h, l)} \theta_{jt}(h, l) \alpha_{jt}(h, l), \quad \text{for } l = 1, \dots, M.$$

The conditions identifying the optimal employment levels for each of these groups of adjusting establishments may then be written recursively as below in the $\sum_{h=1}^M J_h$ equations of (30)-(31). For $l = 1, \dots, M$:

$$\left[D_1 f(n_{0t}(l), a_l, z_t) - w_t \right] + \beta(1 - d) E \left(\frac{\lambda_{t+1}}{\lambda_t} \frac{\Omega_{1, t+1}(l)}{\omega_{0t}(l)} \mid S_t \right) = 0, \quad (30)$$

where, for each $h = 1, \dots, M$,

$$\begin{aligned} \Omega_{jt}(h) &\equiv \sum_{l=1}^M \theta_{jt}(h, l) [1 - \alpha_{jt}(h, l)] \left[D_1 f((1 - d)^j n_{0t}(h), a_l, z_t) - w_t \right] \\ &+ \beta(1 - d) E \left(\frac{\lambda_{t+1}}{\lambda_t} \Omega_{j+1, t+1}(h) \mid S_t \right), \quad \text{for } j = 1, \dots, J_h - 1. \end{aligned} \quad (31)$$

Note that these conditions closely parallel that in the model without plant-level productivity shocks. As there, the marginal effects of a single plant's current employment choice on

its production and wage payments continue for so long as it does not re-adjust its employment. The second term in (30) reflects the probability-weighted sum of future effects for any single member of the $\omega_{0t}(l)$ group of adjusting plants. Any such plant may enter date $t + 1$ with productivity a_1 and produce without readjusting its employment with probability $\frac{\theta_{1,t+1}(l,1)}{\omega_{0t}(l)}[1 - \alpha_{1,t+1}(l, 1)]$; with probability $\frac{\theta_{1,t+1}(l,2)}{\omega_{0t}(l)}[1 - \alpha_{1,t+1}(l, 2)]$, the plant may have productivity a_2 and not adjust employment next period, and so forth. The collections of equations in (31) summarize the resulting marginal effects from $t + 1$ and forward until date $t + J_l$, the date by which this currently adjusting plant will with certainty re-adjust its employment if it has not already done so.

The conditions determining optimal adjustment fractions from within each group of plants are listed below. In each case, the adjusting fraction equates the marginal cost paid to adjust the last plant from a given group to the net value of moving that plant into the adjustment group associated with its current productivity. For each group of establishments that had productivity a_h at the time of last adjustment, $h \in \{1, \dots, M\}$, and now have productivity a_l , $l \in \{1, \dots, M\}$, the optimal adjustment fractions from each time-since-last-adjustment subgroup will satisfy

$$\xi\left(\alpha_{jt}(h, l)\right) = v_{0t}(l) - v_{jt}(h, l) \text{ for } j = 1, \dots, J(h, l) - 1, \quad (32)$$

where $\xi(\alpha) \equiv G^{-1}(\alpha)$. For example, $\alpha_{jt}(1, 2)$, the fraction adjusted from the plant group identified by $(n_{jt}(1), a_2)$, equates the associated marginal adjustment cost to the difference between $v_{0t}(2)$, the value of a plant of type $(n_{0t}(2), a_2)$ at the time of production, and $v_{jt}(1, 2)$, that of a type $(n_{jt}(1), a_2)$ plant at production time. Together, (32) determines all $\sum_{h=1}^M \sum_{l=1}^M [J(h, l) - 1]$ adjustment fractions, given the production-time values of each plant type below.

The values associated with members of each different plant group are expressed recursively as follow. First, the value of any adjusting plant with current productivity a_l is, for

each $l = 1, \dots, M$,

$$v_{0t}(l) = \left[f(n_{0t}(l), a_l, z_t) - w_t n_{0t}(l) \right] + \beta E \left[\frac{\lambda_{t+1}}{\lambda_t} \sum_{m=1}^M \phi(l, m) \left(\alpha_{1,t+1}(l, m) v_{0,t+1}(m) \right. \right. \\ \left. \left. + [1 - \alpha_{1,t+1}(l, m)] v_{1,t+1}(l, m) - \Gamma(\alpha_{1,t+1}(l, m)) \right) \mid S_t \right].$$

This includes the plant's current profit flows associated with having $(n_{0t}(l), a_l)$ at production time, plus its discounted probability-weighted continuation value. At date $t + 1$, the plant will draw productivity a_1 with probability $\phi(l, 1)$. In this case, its expected fixed cost payment will be the conditional average adjustment cost paid by any member of its group, $\Gamma(\alpha_{1,t+1}(l, 1))$, and it will adjust employment to $n_{0,t+1}(1)$ with probability $\alpha_{1,t+1}(l, 1)$, or not adjust with probability $[1 - \alpha_{1,t+1}(l, 1)]$. Alternatively, the plant will have productivity a_2 with probability $\phi(l, 2)$, in which case its expected fixed cost payment will be $\Gamma(\alpha_{1,t+1}(l, 2))$, and it will adjust its employment to $n_{0,t+1}(2)$ with probability $\alpha_{1,t+1}(l, 2)$, or not with probability $[1 - \alpha_{1,t+1}(l, 2)]$, and so forth.

Finally, we have remaining $\sum_{h=1}^M \sum_{l=1}^M [J(h, l) - 1]$ equations that identify the values associated each group of nonadjusting plants in the economy. For each (h, l) past and current productivity combination, for each $j = 1, \dots, J(h, l) - 1$, the value of a nonadjusting plant identified by $(n_{jt}(h), a_l)$ at production time in date t is

$$v_{jt}(h, l) = \left[f(n_{jt}(h), a_l, z_t) - w_t n_{jt}(h) \right] + \beta E \left[\frac{\lambda_{t+1}}{\lambda_t} \sum_{m=1}^M \phi(l, m) \left(\alpha_{j+1,t+1}(h, m) v_{0,t+1}(m) \right. \right. \\ \left. \left. + [1 - \alpha_{j+1,t+1}(h, m)] v_{j+1,t+1}(h, m) - \Gamma(\alpha_{j+1,t+1}(h, m)) \right) \mid S_t \right].$$

Given the current aggregate state, S_t , and the values of $J(h, l)$, the evolution of this economy with persistent plant-specific productivity shocks is fully described by a system of X first-order stochastic difference equations, where

$$X = 7 + 2 \sum_{h=1}^M \sum_{l=1}^M J(h, l) + (M + 2) \sum_{h=1}^M J_h - M^2.$$

Thus, the economy’s aggregate dynamics may be solved as a local appropriation around the steady state using standard linear systems methods.²² In solving for the steady-state, we add to this system the solution for the additional endogenous variables defining the nonadjustment horizons, $\left[[J(h, l)]_{l=1}^M \right]_{h=1}^M$.

7 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 the focus in our numerical examples on cases where idiosyncratic uncertainty at the plant level that is transitory and there are no 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 planning problem.

²²For example, suppose that a took on one of only $M = 2$ possible states, and suppose that the highest possible adjustment cost draw was sufficiently low such that none of the resulting four $J(h, l)$ nonadjustment horizons exceeded 4 periods. Then the maximum size of the resulting linear system would be 163 equations.

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. Moreover, while we have selected to abstract from additional sources of plant-level heterogeneity in the numerical examples here, we have shown how the addition of persistent plant-specific shocks is a straightforward extension to our current framework. 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* 114, 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 [2003a], “Nonconvex factor adjustments in equilibrium business cycle models: do nonlinearities matter?” *Journal of Monetary Economics* 50, 331-360.
- [9] Khan, Aubhik and Julia K. Thomas [2003b], “Inventories and the Business Cycle: An Equilibrium Analysis of (S,s) Policies,” NBER working paper 10078.

- [10] King, Robert G. and Sergio T. Rebelo [1999], “Resuscitating Real Business Cycles,” chapter 14 in M. Woodford and J. Taylor (eds.) *Handbook of Macroeconomics*, vol. IB, Elsevier Science.
- [11] Kollintzas, Tryphon E. [1985], “The Symmetric Linear Rational Expectations Model,” *Econometrica* 53, 963-976.
- [12] Lucas, Robert E. Jr. and Edward C. Prescott [1971], “Investment Under Uncertainty,” *Econometrica* 39, 659-681.
- [13] Mortensen, Dale T. [1973], “Generalized Costs of Adjustment and Dynamic Factor Demand Theory,” *Econometrica* 41, 657-665.
- [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

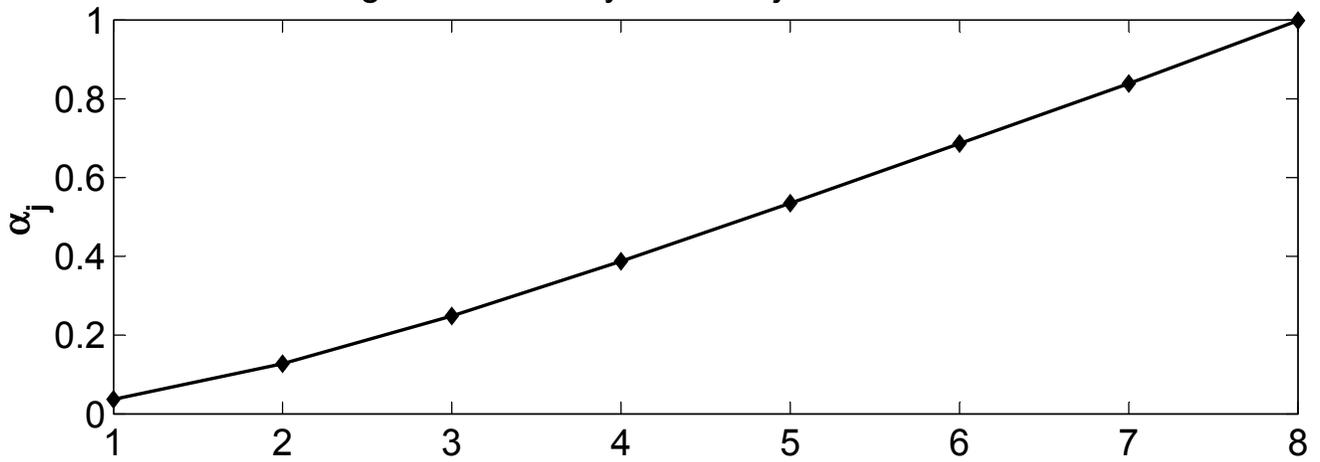


Figure 1b: Steady State Distribution

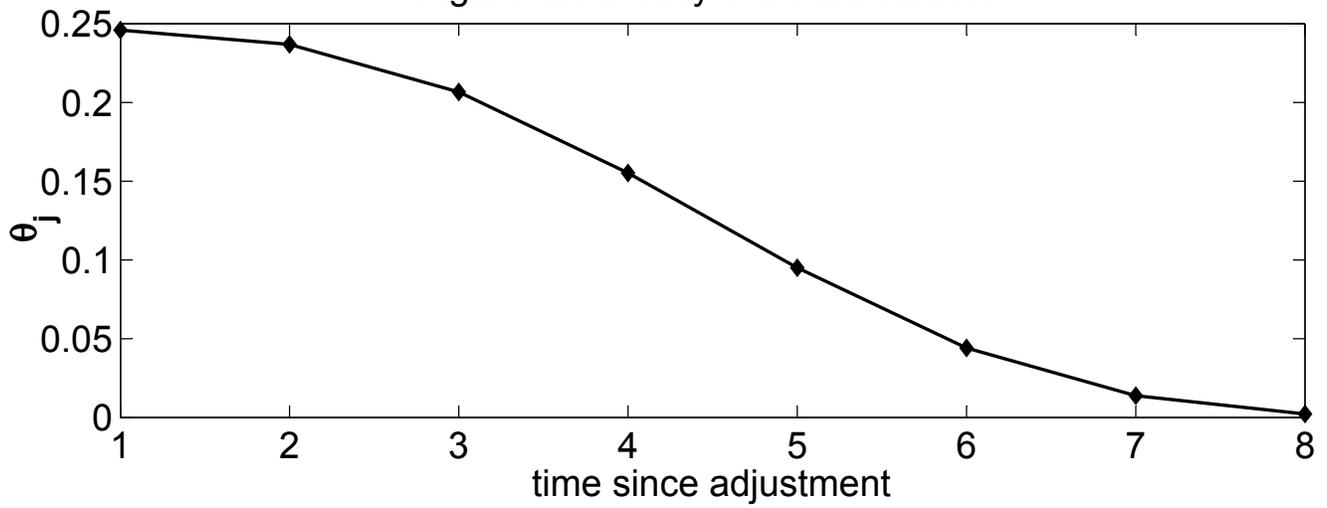


Figure 2a: Market Employment

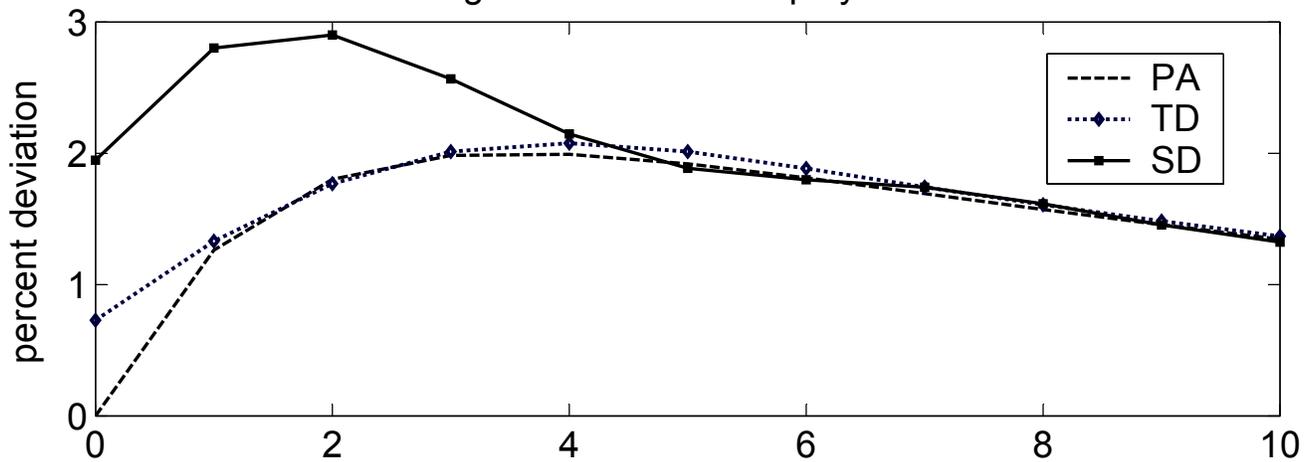


Figure 2b: Market Output

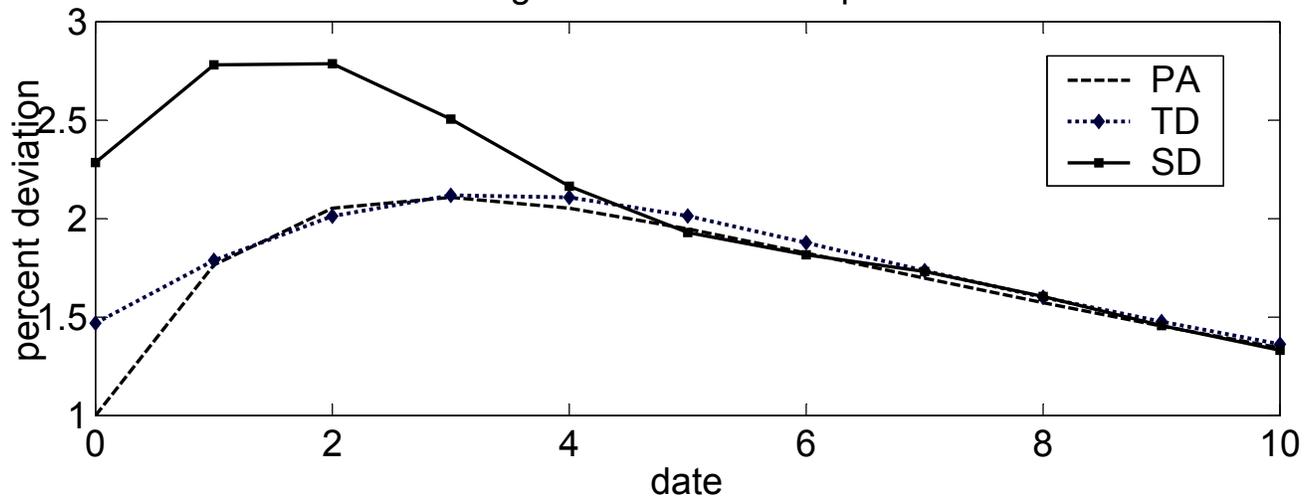


Figure 3a: Total Adjustors

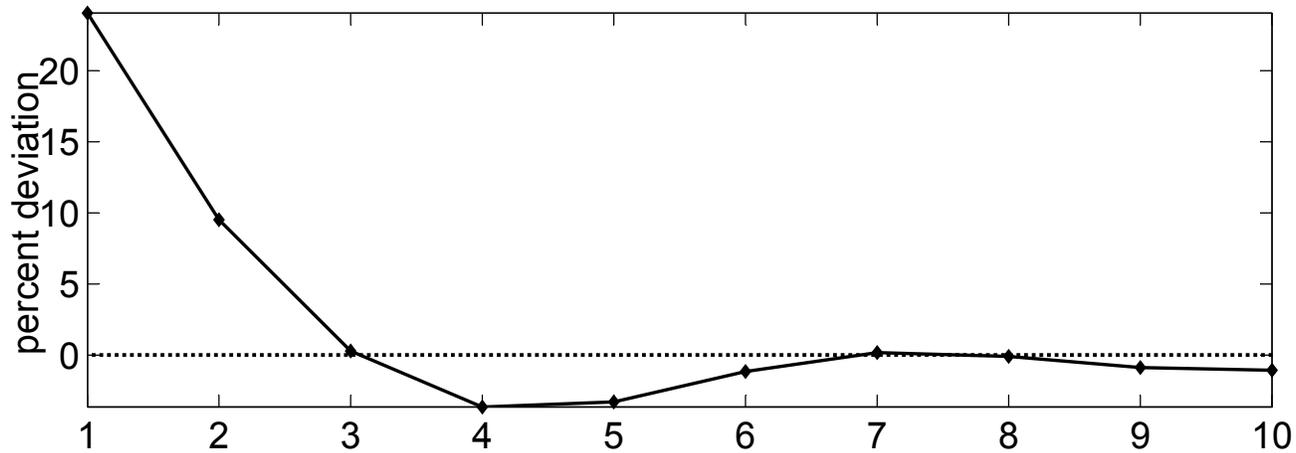


Figure 3b: Target Employment

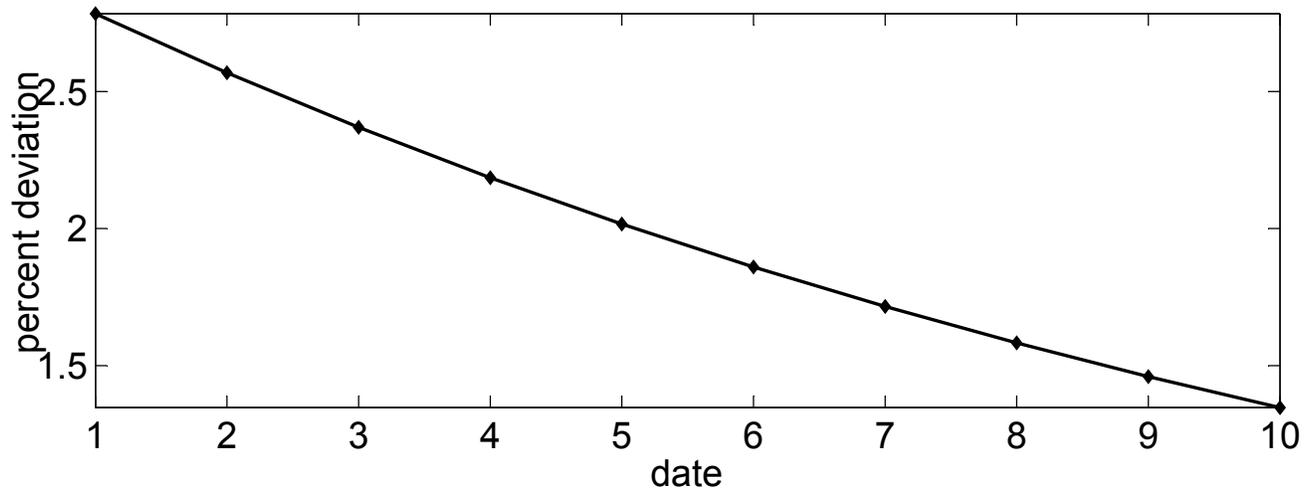


Figure 4: Decomposition of Market Employment

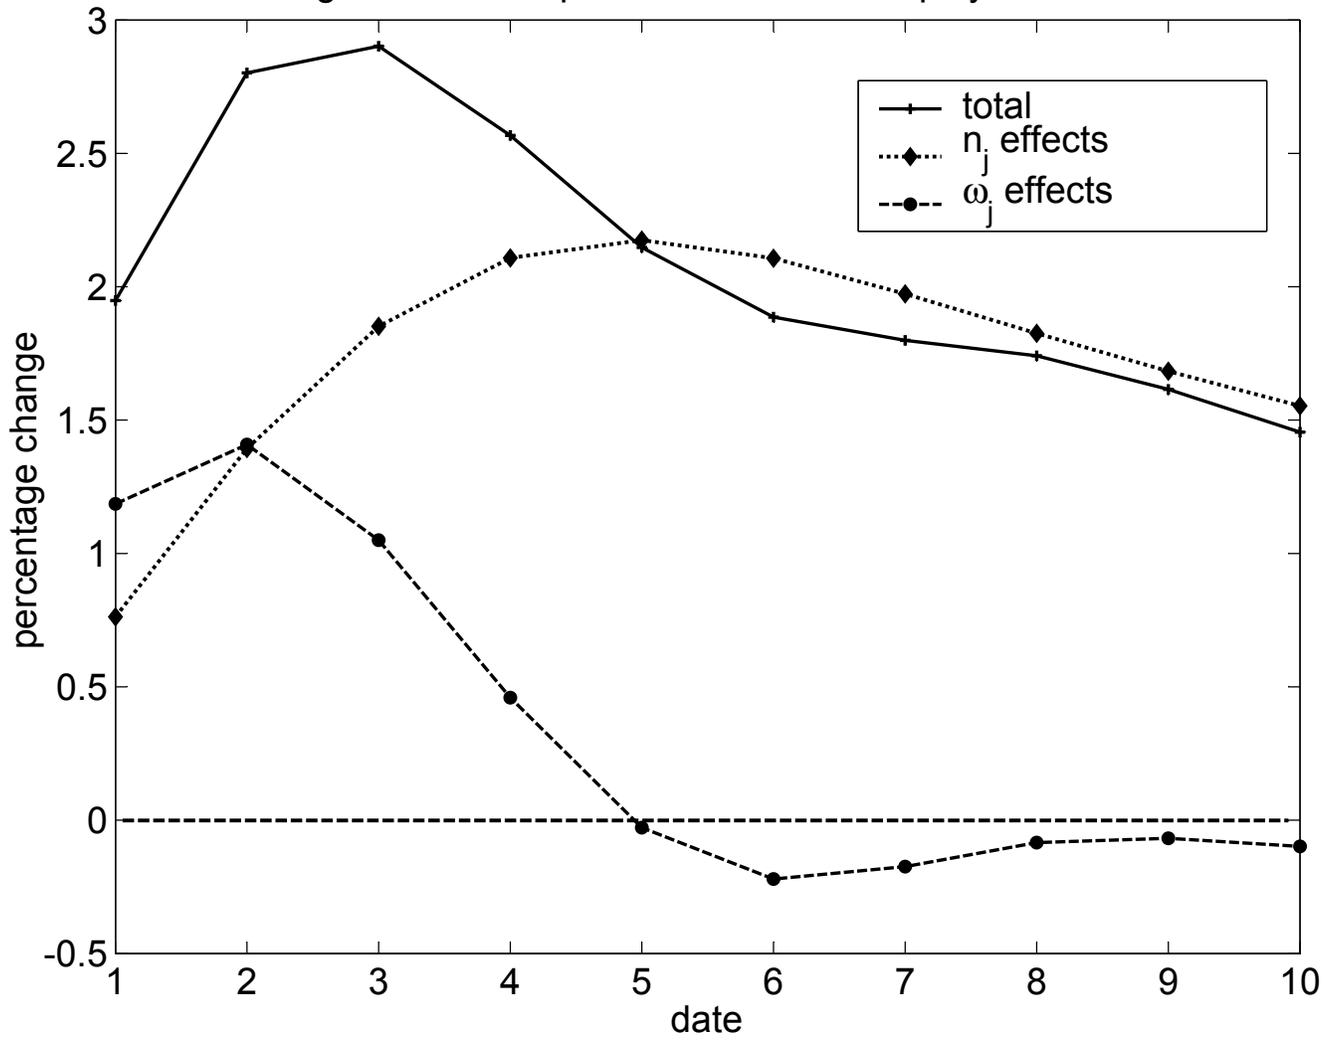


Figure 5: Effects of Equilibrium

