The MIT Press

MICROECONOMIC ADJUSTMENT HAZARDS AND AGGREGATE DYNAMICS*

RICARDO J. CABALLERO AND EDUARDO M. R. A. ENGEL

The basic premise of this paper is that understanding aggregate dynamics requires considering that agents are heterogeneous and that they do not adjust continuously to the shocks they perceive. We provide a general characterization of lumpy behavior at the microeconomic level in terms of an *adjustment-hazard function*, which relates the probability that a unit adjusts to the deviation of its state variable from its moving target. We characterize rich, cross-sectionally dependent *aggregate* dynamics generated by nonconstant hazards. We present an example based on U. S. manufacturing employment and job flows, and find that increasing-hazard models outperform constant-hazard–partial-adjustment models in describing aggregate employment dynamics.

I. INTRODUCTION

Microeconomic units do not adjust continuously to the shocks they perceive, and when they do, adjustments are often large. For example, consumers do not upgrade their cars after every increase in their wealth, and firms do not adjust their factors of production and prices every time demand conditions and productivity change. It is well-known by now that such behavior is inconsistent with the (microeconomic) justification for the standard convex adjustment cost model. This led to the application of (S,s) type models— typically justified by the presence of a nonconvexity in the adjustment cost function—to a wide variety of economic problems wanting for more realistic microfoundations (see, e.g., Barro [1972], Sheshinski and Weiss [1977], Bar-Ilan and Blinder [1987], and Grossman and Laroque [1990]).

Yet, more realism at the microeconomic level does not guarantee more explanatory power at the aggregate level. Intermittent actions are a feature of individual units, not of aggregates. This has led advocates of (S,s) type models to introduce heterogeneity across individual units and to study its aggregate implications. Work in this area goes back to the empirical work on aggregate inventories by Blinder [1981] and is followed by the theoretical steady state

results of Caplin [1985], Caplin and Spulber [1987], and Benabou [1989], and the dynamic (also theoretical) results of Caballero and Engel [1991, 1992a]. All these papers study the case where the incentive for microeconomic adjustment is unidirectional (the "one-sided" (S,s) model).

More recently, work has focused on more general forms of (S,s) models that admit upward and downward adjustments at the microeconomic level. An important virtue of models with this realistic extension is that they can be contrasted with actual data. The work of Bertola and Caballero [1990], and Caballero [1993, 1991], shows that these models can indeed generate *aggregate* dynamics with degrees of persistence and smoothness that are consistent with those observed in reality.[1] Caplin and Leahy [1991] and Caballero and Engel [1992c] show that the long-run average relationships generated by these models are also consistent with several empirical regularities.

The papers in the previous paragraph provide *structural* support for state-dependent models, thereby facilitating the interpretation of the results in terms of "deep parameters." Yet this desirable property is obtained by making assumptions that often limit their empirical flexibility.[2] This motivates the current paper. We provide a "pseudo-structural" framework, where we keep the basic spirit of models in which the decision of adjusting depends on the departure of the main state variable from its target (*state-dependent* models), but do not impose the rigid structure of the simplest (S,s) models; we trade *some* "deep" parameters for empirical richness.

Our approach is far from being a black box, however. We start our description from a given microeconomic policy, without deriving it from first principles. Yet from this point onward, the implications these models have for aggregate dynamics are highly constrained and clearly testable. Furthermore, the hazard functions approach need not be viewed as competing with the structural approach; indeed, it can be part of a sound empirical-theoretical strategy, using it to shed some light on the type of extensions of the structural models that are likely to deliver the highest payoffs.

We consider microeconomic policies that capture what we believe is the main feature of state-dependent models relevant for

---

1. Eberly [1992] and Beaulieu [1991] provide interesting microeconomic evidence validating the microeconomic structure underlying these aggregate models.
2. For example, agents have fixed (S,s) bands over time, which implies that adjustments of the same sign are always of the same size. Furthermore, at the cross-sectional dimension, all agents with adjustments of the same sign adjust by the same amount.

doing empirical work with aggregate series; namely, that the probability that a unit adjusts within a given time period depends on the size of its departure from what would be its choice variable if the frictions it faces were momentarily removed. Furthermore, we argue that it is realistic to consider models in which this probability is eventually *increasing* with respect to the size of the departure.

Our starting point is a generic characterization of discontinuous microeconomic actions in terms of an *adjustment hazard function*. This function determines the probability that a unit adjusts in a given time interval, as a function of its deviation.[3] This approach has important empirical virtues. First, the hazard function may take a wide variety of shapes. Thus, the proposition that it eventually becomes increasing and, simultaneously, that this matters for aggregate dynamics, is testable. Second, in the particular case where the hazard is constant, the model generates aggregate dynamics indistinguishable from those of the quadratic-adjustment-cost–representative-agent model [Rotemberg, 1987]. This establishes a convenient metric to assess the empirical relevance of the state-dependent models discussed here.

The remainder of the paper proceeds as follows. In Section II we introduce the concept of an adjustment hazard function, describe heterogeneity in terms of the endogenous evolution of a cross-sectional distribution of deviations from the microeconomic targets, and obtain an explicit expression for the evolution of the aggregate.

In Section III we characterize the implications of the shape of the adjustment hazard function for aggregate dynamics. Nonconstant hazards, and in particular increasing hazards, are shown to introduce nonlinearities and complex dynamics in aggregate relationships. In contrast with the quadratic-adjustment-cost–representative-agent models, higher moments of the cross-section distribution of deviations have an effect on the aggregate's dynamic behavior for increasing hazard models. First moments are not enough to determine the evolution of the aggregate since higher moments may play an important role in how current innovations are filtered through the cross-section distribution of deviations. These insights are summarized in terms of simple expressions involving higher moments of the cross-sectional distribution of deviations and powers of aggregate shocks.

Section IV illustrates the potential relevance of the results of

3. Adjustment hazard functions are state-dependent. In this sense they differ from the usual hazard functions, where the probability of adjusting depends on how much time has elapsed since adjustment last took place.

the previous sections by providing an example based on the aggregate behavior of net and gross aggregate flows in U. S. manufacturing employment/jobs. The results are encouraging: the estimated hazard function is clearly increasing, and the model outperforms the partial-adjustment model in explaining the dynamic behavior of both net and gross flows. Section V discusses several extensions, and Section VI concludes.

II. ADJUSTMENT HAZARD MODELS AND CROSS-SECTIONAL DISTRIBUTIONS

The basic premise of this paper is that understanding aggregate dynamics requires considering that agents are heterogeneous and that they do not continuously adjust to the shocks they experience. Firms and consumers do not respond frequently to changes in their environment, and when they respond, adjustments are typically large.[4] Possible explanations range from the presence of fixed and proportional costs in the adjustment-cost function to near rationality arguments. Yet our focus in this paper is not on the microeconomic underpinnings of infrequent and lumpy adjustment; instead our aim is to describe the aggregate implications of such a microeconomic environment. In order to facilitate the exposition, and because of the application in Section IV, we use firms' employment/jobs decisions as an *example* to present our framework. Of course, the framework we develop is considerably more general. Among other applications it can be used to study the dynamic behavior of business and residential investment, consumer durables expenditure, inventory investment, and the price level.

Consider a firm $i \in [0,1]$ at time t, that employs e_{it} workers (all variable in logarithms unless otherwise stated) but would employ e_{it}^* workers if frictions were *momentarily* removed, and define the difference between these two quantities (called the *firm's deviation*) as

$$z_{it} \equiv e_{it} - e_{it}^*.$$

We refer to e_{it}^* as the "frictionless" employment level, where the quotation marks are used to stress that in general e_{it}^* does not coincide with the (static) solution obtained when frictions are

---

4. See, e.g., Hamermesh [1989], Davis and Haltiwanger [1990, 1992], and Bresnahan and Ramey [1991] for evidence on lumpy changes in plant level employment.

permanently removed, $\hat{e}_{it}$. We return to this issue in the empirical section.

For expository simplicity we assume that firms set z_{it} equal to zero every time they adjust. Thus, jumps have size $-z_{it}$. This assumption can be relaxed easily to account for the possibility that units may bridge only part of their gap from their "frictionless" level when they decide to jump and to incorporate multiple and stochastic return points. We postpone further discussion of this issue until the final section of the paper.

One of the features that characterizes a particular adjustment-cost model is the rule by which it relates the occurrence and magnitude of adjustments to the size of firms' departures. For example, when adjustment costs are quadratic, the firm adjusts continuously and proportionally to z_{it}. Alternatively, in the case of nonconvex adjustment cost (S,s) models, adjustment takes place only when z_{it} reaches certain thresholds.

In this paper we capture what we believe is the most distinguishing feature of state-dependent models by describing the (discontinuous) microeconomic adjustment policy in terms of an "adjustment hazard rate function" (adjustment hazard or hazard function for short). We assume that the probability that firm i adjusts its level of employment during the (small) time interval $(t, t + dt)$ is (approximately) equal to $\Lambda(z_{it})dt$, where $\Lambda(z)$ denotes the adjustment hazard. The disequilibrium variable z_{it} determines how likely it is that a firm adjusts its level of employment in a given time period.

The adjustment hazard framework is quite general. For example, it includes the family of (S,s) models as the particular case where the adjustment hazard function is equal to zero within the inaction range and infinity elsewhere. It also includes the representative-agent–quadratic-adjustment-cost model (or partial-adjustment model) as a particular case, since the aggregate dynamics of this model are indistinguishable from those of the constant hazard model [Rotemberg, 1987].

In principle, a hazard function could take almost any shape. In practice, however, reasonable hazard functions should eventually be (strongly) increasing with respect to the absolute value of z_{it}, since it is improbable that firms tolerate large departures as well as they tolerate small departures. We call this realistic feature the *increasing hazard property,* and study its implications for aggregate dynamics in detail. We note that the family of (S,s) models corresponds to an extreme case of increasing-hazard models, where

the hazard function jumps from zero to infinity at the trigger points.

In an adjustment-hazard model, the dynamics of employment are determined by the interaction between the shape of the hazard function and the shifts in the cross-sectional density induced by aggregate and idiosyncratic shocks. The cross-sectional density, denoted by $f_t(z)$, plays a central role in the approach we present here (as it does in most adjustment-cost models); the effect of the same sequence of shocks typically depends on the initial cross-sectional distribution of firms' deviations. Furthermore, precisely because of this dependence, the response of aggregate employment to aggregate shocks is generally nonlinear and exhibits complex dynamics. For example, if the hazard function is increasing and history is such that most firms' deviations are small (i.e., the absolute value of the z_i's is small), then the number of firms responding to an aggregate shock is small. On the other hand, if most z_i's are large (in absolute value), then this number is large.

Formally, and working in discrete time to simplify the exposition, the change in aggregate labor demand during the time interval $(t, t + 1]$ is equal to[5]

$$(1) \qquad \Delta E_{t+1} = \int_{-\infty}^{\infty} (\Delta E_{t+1}^* - z) \Lambda(z - \Delta E_{t+1}^*) f_t(z) \, dz,$$

where E and E^* denote aggregate employment and "frictionless" employment, respectively.[6] To derive this equation, we first consider the fraction $f_t(z)$ of firms with deviation z at time t, just after firms have experienced their idiosyncratic shocks. After an aggregate shock that leads to a change in aggregate "frictionless" employment of ΔE_{t+1}^* takes place, these firms' z_i's change to $z - \Delta E_{t+1}^*$; a fraction $\Lambda(z - \Delta E_{t+1}^*)$ among them adjust with the hazard shock, all of them by $(\Delta E_{t+1}^* - z)$. Adding over all possible values of z yields (1).

The evolution of the cross-sectional distribution between time periods t and $t + 1$ is determined by the aggregate shock, the firms that adjust, and the new idiosyncratic shocks. The cross-sectional density at time $t + 1$ then summarizes the relevant history for the next sequence of shocks. Formally, we define an operator T_{t+1} that

5. The assumption that the number of firms is large (a continuum) is implicit here, since we assume that among all firms that have deviation z just before the hazard shock, the fraction that adjusts is equal to $\Lambda(z)$.
6. We assume that all establishments face the same hazard function, and we discuss the case with heterogeneous hazards in the final section. A more general expression for ΔE_{t+1} which allows for the possibility of idiosyncratic shocks (beyond the hazard shock) can be derived analogously.

maps the cross-sectional density at time t, f_t, and the change in aggregate "frictionless" employment, $\Delta E_{t+1}^* \equiv E_{t+1}^* - E_t^*$, into the new cross-sectional density:

$$(2) \qquad f_{t+1} = T_{t+1}(f_t, \Delta E_{t+1}^*),$$

where the operator T depends not only on the hazard function $\Lambda(z)$ and the size of the aggregate employment shock, ΔE_{t+1}^*, but also on the stochastic mechanism underlying idiosyncratic (firm-specific) shocks (beyond the hazard shock). At this point we do not need to be more precise about this operator; later, in the application section we provide a simple example of it.

III. AGGREGATE DYNAMICS OF ADJUSTMENT-HAZARD MODELS

In the previous section we presented the basic elements required to track down aggregate dynamics when microeconomic units adjust intermittently to the shocks they perceive. In this section we concentrate on the implications of nonconstant hazard models for aggregate dynamics and describe in detail how the microeconomic hazard function interacts with the evolution of the cross-sectional distribution, thereby determining the response of aggregate employment to aggregate shocks. We first consider three particular families of hazard functions and then derive general expressions for employment dynamics.

III.1. Constant Hazard

It is instructive to begin by describing the constant-hazard model [Calvo, 1983]. This case generates aggregate dynamics identical to the *linear* dynamics of the partial-adjustment model, which in turn can be obtained from a representative-agent framework with quadratic adjustment costs.[7] Since the latter is a specification often used by macroeconomists to characterize aggregate dynamics, it constitutes a convenient benchmark for a discussion of the more realistic increasing-hazard models.

---

7. See Rotemberg [1987] for a proof of the aggregate equivalence between constant hazard models and representative agent models with quadratic adjustment costs. What this result says is that, based on an aggregate employment series, it is impossible to distinguish between an economy where all firms adjust their employment levels in every period by a fraction λ of their current deviation, from an economy where a fraction λ of agents adjusts completely in every period. In the final section of this paper we argue that this lack of identifiability holds only for the constant hazard case. It is important to note that Rotemberg's [1987] equivalence result goes beyond the points stressed here, since it incorporates the optimal determination of e_{it}^*.

Given the simple structure of the adjustment-hazard function we consider in this example, ($\Lambda(z) \equiv \lambda_0$), equation (1) simplifies to

$$(3) \qquad \Delta E_{t+1} = \lambda_0(\Delta E^*_{t+1} - Z_t^{(1)}),$$

where $Z^{(k)}$ denotes the kth moment of the cross-sectional distribution of deviations. This equation shows that the dynamics of the constant hazard model depend on the cross-sectional distribution only through its mean. Since the adjustment hazard is constant, the *fraction* of firms that adjust their employment levels in every period does not change over time; it is only the average *magnitude* of these adjustments that depends on the sequence of previous shocks.

If $Z_t^{(1)}$ is relatively low—i.e., there is a relatively large share of upward adjustments in firms' employment levels that has not occurred—then ΔE_{t+1} will be larger for all values of the realization of the aggregate shock, and the opposite will happen if $Z_t^{(1)}$ is relatively high. Yet, this example does not generate nonlinear responses to aggregate shocks, nor does it have a role for "distributional" effects (in the usual higher moments sense). An expression showing this can be derived by replacing the definition $Z_t^{(1)} \equiv E_t - E^*_t$ in equation (3), which yields

$$\Delta E_{t+1} = (1 - \lambda_0)\Delta E_t + \lambda_0 \Delta E^*_{t+1}.$$

Thus, the dynamics generated by a constant-hazard model are straightforward and can be captured entirely by a simple first-order autoregressive term; this is not the case for models with nonconstant-adjustment hazards.

III.2. Simple Asymmetric-Hazard Model

Before describing the dynamics of an increasing-hazard model—which, as argued above, provides a more realistic description of actual microeconomic adjustments—it is useful to illustrate the emergence of "distributional effects" with one of the most basic departures from the constant-hazard model: the simplest piecewise constant hazard, where the probability of adjusting depends on whether the firm's deviation is positive or negative. We therefore consider the following asymmetric-hazard function:

$$\Lambda(z) = \begin{cases} \lambda^+, & \text{for } z > 0, \\ \lambda^-, & \text{for } z < 0. \end{cases}$$

Replacing this hazard in equation (1) yields the following expres-

sion for aggregate changes in employment:

$$(4) \qquad \Delta E_{t+1} = \lambda^- F_t(\Delta E^*_{t+1})[\Delta E^*_{t+1} - Z_t^{(1-)}(\Delta E^*_{t+1})]$$
$$+ \lambda^+(1 - F_t(\Delta E^*_{t+1}))[\Delta E^*_{t+1} - Z_t^{(1+)}(\Delta E^*_{t+1})],$$

where $F_t(z)$ denotes the cross-sectional distribution function at time t evaluated at z, and $Z_t^{(1-)}$ (ΔE^*_{t+1}) and $Z_t^{(1+)}$ (ΔE^*_{t+1}) the means of the cross-sectional distribution conditional on z being smaller and larger than ΔE^*_{t+1}, respectively.

Even in this simple case, the response of aggregate employment to (contemporaneous) aggregate shocks is nonlinear. For example, consider the case where the probability of firing workers during a given time period is larger than that of hiring workers ($\lambda^+ > \lambda^-$), and suppose that an aggregate shock that increases the "frictionless" levels of employment takes place. This shock shifts the cross-sectional density to the left. Thus, even though the size of a firm's jump (if it should adjust) changes by ΔE^*_{t+1}, these jumps become less likely since the fraction of firms facing a small probability of adjusting is larger. The situation is reversed when we consider the effect of a negative aggregate shock. The exact form of the relation between the change in aggregate employment and the aggregate shock depends crucially on higher moments of the cross-sectional density before the shock.

III.3. Simple Increasing-Hazard Model

We illustrate the main aggregate features of an increasing-hazard model through a simple quadratic case:

$$\Lambda(z) = \lambda_0 + \lambda_2 z^2,$$

with $\lambda_2 > 0$.[8] We use (1) to find an expression for aggregate employment changes:

$$(5) \qquad \Delta E_{t+1} = \lambda_0 \Delta E^{*pa}_{t+1} + \lambda_2\{(\Delta E^{*pa}_{t+1})^3 + 3\Delta E^{*pa}_{t+1} Z^{(2)}_{c,t} - Z^{(3)}_{c,t}\},$$

where $\Delta E^{*pa}_{t+1} \equiv (\Delta E^*_{t+1} - Z_t^{(1)})$ corresponds to the partial-adjustment term in equation (3). We dub this expression the *effective aggregate shock*; it is positive when the contemporaneous shock is large relative to the unadjusted portion of previous shocks and negative otherwise. The terms $Z^{(k)}$ and $Z_e^{(k)}$ denote the kth moment (noncentered and centered, respectively).

The expression in brackets in equation (5) has several interesting features. First, it is nonlinear with respect to contemporaneous

8. Since the time period considered when sampling has length one, we have that, strictly speaking, the hazard function is equal to $\max(\min(\Lambda(z),1),0)$.

aggregate effective shocks. This nonlinearity grows with the size of the shock. Second, aggregate shocks interact with the moments of the cross-sectional distribution: a larger variance leads to larger responses of aggregate employment to aggregate shocks. The corresponding term has the same sign as the effective shock and a magnitude proportional to the variance of the cross-sectional distribution of deviations. Finally, there is a pure "history" term (beyond that present in the partial-adjustment model). For a given effective shock, increases in employment are larger the more skewed to the left the distribution of deviations is.

We present a simple illustrative example: the initial cross-sectional density is symmetric with respect to the target level, so that the effective and actual aggregate shock are the same. At this point an aggregate shock (of size ΔE^*) increases every firm's "frictionless" employment level by an amount equal to the size of the shock (i.e., all z_i's decrease by ΔE^*). Following this shock, the *hazard shock* takes place, whereby some firms decide to change their employment levels, while others remain inactive. Since the hazard function is increasing, the fraction of firms that decides to hire workers is larger when the shock is large. Therefore, the number of workers hired increases more than one for one with the size of the shock. Similarly, the fraction of firms that fires workers decreases, and the number of workers fired decreases more than one for one. It follows that the net effect of an aggregate shock on employment is nonlinear in the size of the shock. Figure I shows the initial cross-sectional density (solid line), this density after a small shock (dashed line), and the same density after a shock that is twice as large (dashed-dotted line). An increasing-hazard function is also included in this figure. Since the change in employment is equal to the weighted average of all possible jumps, it follows from this figure that the increase after the large shock is more than twice as large as the corresponding increase after the small shock. Figure II illustrates the effect of second moments. The solid and dashed lines in the figure show two densities with the same first moment but different second moments; a quadratic hazard function is also included. We learn from the figure that the distribution with more weight on the tails has a higher average hazard, and therefore describes a situation where aggregate employment is more responsive to aggregate shocks.

III.4. The General Case

It is apparent by now that, within the class of discontinuous microeconomic actions models studied here, the absence of distribu-

FIGURE I
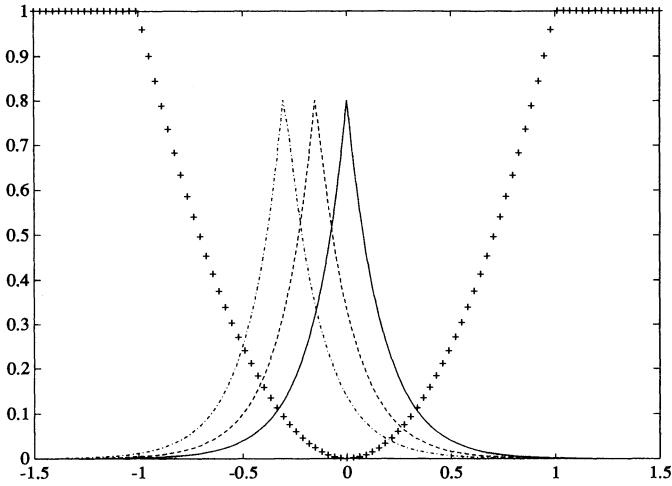The Cross-Sectional Density Before and After a Small and Large Shock

tional effects and nonlinearities in aggregate dynamics is a special feature of the partial-adjustment model—a model that is highly unlikely to be a good description of microeconomic behavior. In a sense, this has always been known, or at least suspected. Our aim
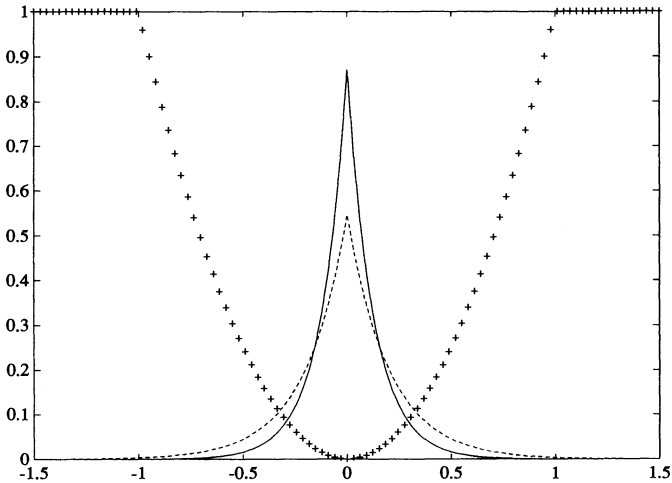

FIGURE II
The Effect of Aggregate Shocks Depends on Higher Moments

in this paper is to provide a concrete formalization of these ideas and a tractable methodology to integrate these realistic effects into standard empirical work. The following representation of aggregate dynamics under a general hazard function is useful for this purpose:

$$(6) \qquad \Delta E_{t+1} = \sum_{k \geq 0} a_k (\Delta E_{t+1}^{*pa}) Z_{c,t}^{(k)},$$

where the $a_k(\cdot)$ functions are obtained by using a Taylor expansion of the hazard function around $z = \Delta E_{t+1}^{*pa}$ in (1).[9] Hence,

$$a_0 (\Delta E_{t+1}^{*pa}) = \Lambda (\Delta E_{t+1}^{*pa}) \Delta E_{t+1}^{*pa},$$

and

$$a_k (\Delta E_{t+1}^{*pa}) \equiv (-1)^k \left[\Delta E_{t+1}^{*pa} \frac{\Lambda_k (\Delta E_{t+1}^{*pa})}{k!} + \frac{\Lambda_{k-1} (\Delta E_{t+1}^{*pa})}{(k-1)!} \right],$$

where Λ_k denotes the kth derivative of the adjustment hazard function, and $k \geq 1$.[10]

Assuming that the probability of adjusting employment depends *only* on a firm's deviation, as we have done so far, ignores many external features that may influence firms' decisions, such as changes in adjustment costs due to strategic interactions and congestion effects. These issues can be incorporated into the adjustment-hazard approach by considering hazard functions of the form $\Lambda(z,x(t))$, where $x(t)$ denotes a set of variables that affect the probability of adjusting employment *conditional on the size of the deviation*. Yet in the absence of a theoretical justification for the particular family of hazard functions being used, we should be conservative when interpreting time-varying hazard functions. With this caveat, time-varying hazards may provide a clue as to whether time-varying structural (S,s) bands are likely to be empirically relevant.

IV. EXAMPLE

In this section we apply the methodology described in the previous section to postwar U. S. manufacturing employment

---

9. See subsection III.3 for the definition and economic interpretation of ΔE_{t+1}^{*pa}.
10. Expressions similar to the one obtained above can be derived for *gross* flows (job creation and job destruction) too.

data.[11] We show how in this case increasing-adjustment-hazard models outperform partial-adjustment models in fitting both net and gross employment/job flows data.

IV.1. "Frictionless" Model

The first stage in estimating a model of the type described here is to characterize the behavior of a firm's "frictionless" component, Δe_{it}^*. This is then used to construct an estimate of the path of its aggregate counterpart, ΔE_t^*, which is the basic input (the aggregate shocks) for the problem at hand.

We let the ith firm's production function and demand at time t be given by

$$y_{it} = (\alpha e_{it} + \beta h_{it}) + \epsilon_{it},$$
$$p_{it} = -(1/\eta)y_{it} + v_{it},$$

where y_{it}, h_{it}, ϵ_{it}, p_{it}, and v_{it} denote output, hours per worker, productivity level, price and demand shock, respectively. The parameter η is the price elasticity of demand faced by firm i. We follow Bils [1987] and assume that $\beta > \alpha \approx 1$. We also assume that the firm is competitive in the labor market, but faces a (per hour) wage curve that is a function of the average number of hours worked: $w_{it} = g(h_{it}) + w_t$. The functional forms we have chosen imply that, in the absence of employment adjustment costs, the firm always chooses the same number of hours worked per worker, and adjusts to productivity and demand shocks only by varying employment.

Under the assumption that increments in productivity and demand levels are approximately independent (over time for each i), we can approximate the e_{it}^*'s—up to an additive constant—by the corresponding static frictionless values, $\hat{e}_{it}$. That is, we let the firm maximize current revenues with respect to employment— facing no adjustment costs—given that h_{it} is at its frictionless

---

11. All employment data are quarterly and seasonally adjusted, and cover only manufacturing production workers. We use two measures of net employment flows. The first one is constructed by the Bureau of Labor Statistics, and the second one corresponds to series constructed by Davis and Haltiwanger [1990, 1992] from the LRD. The correspondence between these two sources is extremely high at the level of aggregation we use in this paper, both for the employment levels and rates of growth. Gross flows are those in Davis and Haltiwanger and correspond to the sum of net changes at the establishment level. That is, job creation is the sum of employment changes in all those establishments that had a positive change in employment during the quarter, while job destruction corresponds to the sum of the negative net changes. Hours correspond to the BLS measure of *average* hours worked. The sample periods used vary and are discussed when presenting the results.

optimal value $\bar{h}$. The first difference of the resulting expression can be used to describe Δe^* in terms of current (unobserved) demand and productivity shocks:

$$(7) \qquad \Delta e_{it}^* = (\gamma\Delta\epsilon_{it} + \Delta v_{it} - \Delta w_t)/(1 - \alpha\gamma)$$

where $\gamma \equiv (\eta - 1)/\eta$ and, to satisfy the second-order conditions, $\alpha\gamma < 1$.

To obtain an expression for Δe^* in terms of observables, we assume that there are no adjustment costs in average hours (for evidence on this see, e.g., Sargent [1978] and Shapiro [1986]). Thus, firms choose hours optimally in response to productivity and demand shocks, even in the short run. This yields

$$(8) \qquad \Delta e_{it} = [\gamma\Delta\epsilon_{it} + \Delta v_{it} - \Delta w_t + (\beta\gamma - \mu)\Delta h_{it}]/(1 - \alpha\gamma)$$

where $(\mu - 1)$ is the elasticity of the marginal wage schedule with respect to average hours worked.[12] It is now straightforward to factor out demand and productivity shocks, as well as the time-dependent component of the wage schedule, yielding

$$(9) \qquad \Delta e_{it}^* = \Delta e_{it} + \theta\Delta h_{it},$$

where $\theta \equiv (\mu - \beta\gamma)/(1 - \alpha\gamma)$. Aggregating equation (9) over all firms yields an expression that can be used to construct an estimate of the path of aggregate "frictionless" labor demand:[13]

$$(10) \qquad \Delta E_t^* = \Delta E_t + \theta\Delta H_t.$$

If productivity shocks are smooth, we can obtain an alternative expression for ΔE_t^* by combining the production function and equation (10):

$$(11) \qquad \Delta E_t^* = c_0 + \frac{1}{\alpha}\Delta Y_t + \left(\theta - \frac{\beta}{\alpha}\right)\Delta H_t,$$

with c_0 a constant. We use the latter expression in Caballero and Engel [1992b] and find hazard functions similar to the ones estimated below, which are based on equation (10).

---

12. When the component of the wage equation that depends on average hours is of the form $G(\mathscr{H}_{it}) = c_1\mathscr{H}_{it}^\mu + c_2$, with $\mathscr{H}$ denoting actual average hours worked instead of logs, this elasticity is constant, and the expression above therefore involves no approximation.

13. We assume that firms' shares in employment, frictionless employment, and average hours worked per worker are independent of idiosyncratic and aggregate shocks.

From Bils [1987] we obtain a value of μ around 1.9.[14] Based on Bils [1987] and Shapiro [1986], we set $\alpha = 1$ and $\beta = 1.06$. Combining these with a markup value of approximately 25 percent yields $\theta = 5$, which is our base case. The qualitative conclusions obtained when comparing increasing hazard models with the partial-adjustment model are unaffected by large variations in θ, however.[15]

IV.2. Estimation Methodology

Estimating the best hazard model within a particular parametric family requires choosing a criterion by which the performance of different adjustment hazards can be compared. A natural candidate is to calculate the series of *net* flows in employment determined by a particular set of parameters, and then look at the sum of the corresponding squared residuals.[16]

We work in discrete space and time. Firms' deviations are allowed to take one of 99 equally spaced values between -1.5 and 1.5; time evolves in quarters. We generate the sequence of cross-section densities as follows: the cross-section density at time $t + 1$ is obtained from that at time t by first shifting the latter by an amount equal to (the negative of) the aggregate shock, ΔE^*_{t+1}, then applying an idiosyncratic shock of size σ_I, so that half the firms with deviation z have their deviation increase by σ_I and the other half have it decrease by σ_I,[17] and finally applying the hazard shock (so that the probability density at a point $z \neq 0$ decreases by a fraction $\Lambda(z) \cdot dt = \Lambda(z)$).[18] This, in conjunction with equations (1),

14. For this we evaluate the elasticity of Bils's marginal wage curve (method 1 in his paper) with respect to hours at 40 hours per week and take 1980 as the base year.

15. The parameter θ cannot be estimated directly using the procedure described below because a value of θ equal to zero yields a perfect fit.

16. When we estimate the hazard function, there is a certain sense in which the variable to be explained, ΔE_t, is a "right-hand-side" variable too (see equation (10)). To ensure that this is not what accounts for the quality of the fit we obtain, we reestimated the models considered later in this section substituting ΔH_t in (9) by a simulated series with the same moments: the corresponding sum of squared residuals (SSR) was many times larger in every case. For example, when compared with the sum of squared residuals obtained with BLS data in Table I, using "noise" instead of changes in hours worked increases the SSR by a factor of 4.8.

17. This discretization scheme differs from the standard discretization for Brownian motion in that it does not impose a relation between how the state space and time are discretized [Engel, 1991, Ch. 3]. With the standard discretization the deviations take values that depend on σ_I, thereby making the estimation process more difficult.

18. This description provides an explicit example of the operator T_t in equation (2). We assume that there is only one shock per quarter, which implies that $\Lambda(z) \leq 1$. Also, since the z-space we work in is bounded, we are implicitly assuming that for values of z beyond the range considered we have that $\Lambda(z) = 1$.

(2), (10), and (11), generates an estimate of the path of ΔE.[19] Since the model derived in subsection IV.1 does not incorporate quits, a 3 percent quarterly quit rate was added exogenously. This is an approximate measure of the drift in the path of the individual z_i's, since costly decisions may be related to jobs rather than to specific workers. We also disregard the cyclical behavior of quits.

There are many macroeconomic problems (e.g., wage determination in search-bargaining models) where it is important to know not only the *net* but also the *gross* employment/job flows (creation and destruction). Since these flows are a natural corollary of the nonrepresentative agent models we consider in this paper, we also incorporate *aggregate* gross flows in estimation (job creation and destruction, in Davis and Haltiwanger's [1992] terminology). In this case we choose the model (within a given parametric family) that minimizes the sum of the norms of the vectors with components equal to the estimated creation and destruction errors. Weights are determined by the inverse of the covariance matrix we obtain in the case with net flows.[20]

IV.3. Estimation: Net Flows

We first estimate the partial-adjustment model (P.A.M.), $\Lambda(z) \equiv \lambda_0$, and quadratic hazard model, $\Lambda(z) = \lambda_0 + \lambda_2(z - z_0)^2$, for the (net) rate of change in U. S. manufacturing employment during

---

19. The initial cross-section density is assumed to be equal to the ergodic density that would exist if there were no aggregate surprises and idiosyncratic shocks followed a random walk with drift equal to that of the aggregate shock process and (instantaneous) variance equal to the sum of the idiosyncratic variance and the variance of the series of aggregate shocks. This is the best choice of initial density in a precise sense (see Caballero and Engel [1992c]). Since it is not the actual initial density, we discard the first four observations when calculating the sum of squared residuals. To check that choosing the initial density in this way has little impact on the estimates we obtain, we started off 44 quarters before the period of interest and applied the corresponding aggregate shocks to the initial ergodic distribution (with BLS data). The effect of the choice of initial distribution clearly washes away in this case, and the cross-sectional distribution at the beginning of the time period of interest therefore is a good approximation of the "true" distribution. The parameter estimates we obtained were similar to the ones obtained with the less time-consuming approach described above.

The fact that ΔE^*_{t+1} and σ_I (typically) differ from integer multiples of the basic step in z-space, h, was taken into account by considering a weighted average of the cross-section distributions that attains when jumps are equal to the two integer multiples of h nearest to the actual jump.

20. The estimated variances of the implied errors in the creation and destruction series are 0.0110 and 0.0115, respectively. The correlation between both error series is 0.557. These parameters were obtained with the BLS data. As a robustness check we also minimized a weighted average of the sums of squared residuals of the creation and destruction series, with weights inversely proportional to the relative variances. The resulting estimates did not differ significantly.

TABLE I
NET FLOWS

Λ	BLS		D-H	
	P.A.M.	Quadratic	P.A.M.	Quadratic
λ_0	0.229	0.019	0.236	0.000
	(0.021)	(0.010)	(0.023)	(0.011)
λ_2	—	0.530	—	0.538
		(0.010)		(0.013)
z_0	—	−0.816	—	−0.840
		(0.011)		(0.028)
σ_I	—	0.059	—	0.077
		(0.015)	—	(0.016)
SSR	0.00555	0.00276	0.00668	0.00445

Notes. Standard errors are in parentheses. SSR: sum of squared residuals. P.A.M.: partial-adjustment model (σ_I not identified in this case). Quadratic: quadratic hazard model (with the parameterization discussed in the text).

the period 1972:1–1986:4.[21] The first column in Table I presents the results for the partial-adjustment model; the estimated constant hazard parameter is 0.229 and significant. The second column contains the quadratic hazard model results. The resulting hazard function is clearly asymmetric and eventually increasing; it is depicted by the solid line in Figure III.[22] Most interestingly, using a quadratic hazard model instead of a partial-adjustment model increases the R^2 coefficient from 0.75 to 0.88; the sum of squared residual of the partial-adjustment model is twice as large as that of the quadratic hazard model.

It may seem as if the set of deviations implied by the model is too large. However, Davis and Haltiwanger [1992] as well as Bresnahan and Ramey [1991] show that it is not rare to find very large adjustments in employment levels. Moreover, from the corresponding ergodic distribution we conclude that a given firm's

21. We choose this period for comparability with gross flow data. The particular parameterization we use for a second-degree polynomial is motivated by the fact that the three parameters are easy to interpret: the hazard function attains its extreme value at $z = z_0$; and this value is equal to λ_0. Finally, λ_2 captures the curvature of the parabola, and its sign determines whether the parabola is convex or concave.

22. We provide a plausible explanation for the asymmetry, and the fact that the hazard function does not attain its minimum near zero, in Section V.
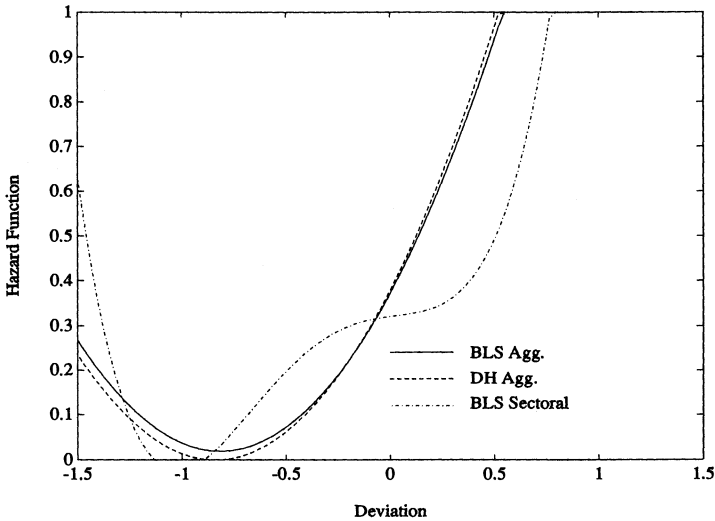
FIGURE III
Different Estimates of the Hazard Function

deviation belongs in a much smaller fixed range of width 0.60 approximately 80 percent of the time.

Figure IV plots the residuals of both models, together with a linear transformation of changes in employment that serves the role of a business cycle clock. This figure shows that a substantial fraction of the gain in terms of fit yielded by the increasing hazard model comes during sharp contractions and brisk expansions.[23] For example, consider the recession of 1974–1975 followed by the expansion of 1975–1976. The relatively large response of employment to aggregate shocks observed during this period (not shown in the figure), requires a flexible short-run elasticity of employment with respect to these shocks. As discussed in subsection III.3 (see Figure I and equation (6)), an increasing-hazard model has this property because, in contrast with the partial-adjustment model, the number of units adjusting varies as the cross-sectional density moves and interacts with different regions of the hazard function.

Both for better comparability with the gross flows results presented below and as a check of robustness, we repeat the

23. Caballero [1992] shows, in the context of (S,s) models, that the probabilistic mechanisms underlying state-dependent models tend to offset the aggregate effects of nonlinearities at the microeconomic level. It follows that these nonlinearities are likely to permeate aggregate dynamics only when shocks are large.
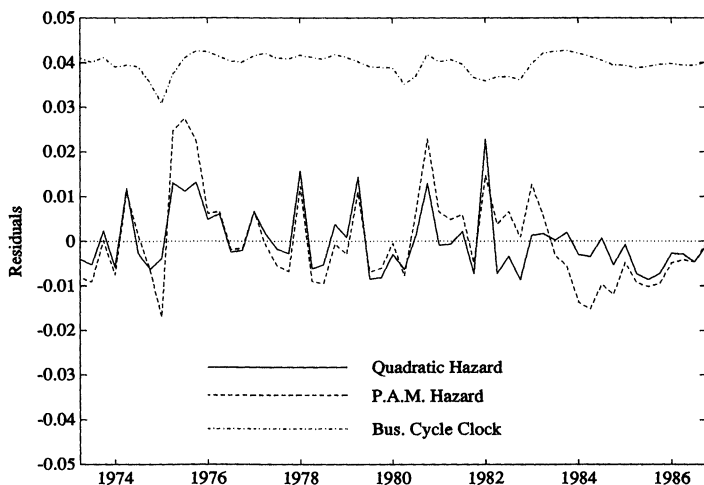
FIGURE IV
Residuals for Increasing and Constant Hazard Models

procedure with Davis and Haltiwanger's [1992] job flow series. Using their net changes in jobs' series, we generate the results in columns 3 and 4 of Table I. Although less dramatic in terms of the SSR gain, the results are similar to those in the first two columns. Moreover, the estimated increasing hazard (short dashes in Figure III) is almost indistinguishable from that estimated with the BLS employment series (solid line in Figure III).

IV.4. Estimation: Gross Flows

Economies are characterized by large amounts of heterogeneity. This is particularly true for creation and destruction of jobs across U. S. manufacturing plants (see, e.g., Davis and Haltiwanger [1990, 1992]). Heterogeneity not only leads to large and simultaneous constant creation and destruction flows that are appended to the cyclical variations in employment, but also gives somewhat independent life to both margins. The correlation between the Davis-Haltiwanger [1990, 1992] creation and destruction flows is only -0.13, and destruction is substantially more cyclical than creation. In this section we study the extent to which a nonconstant-hazard model can capture the rich behavior of gross flows and, more importantly, whether the estimated hazard function required to do so is consistent with the hazard required to explain net flows.

Table II presents the results. Comparing the two columns

TABLE II
GROSS FLOWS

Λ	P.A.M.	Quadratic
λ_0	0.224	0.025
	(0.024)	(0.014)
λ_2	—	0.278
		(0.131)
z_0	—	−1.068
		(0.223)
σ_I	—	0.103
		(0.017)
SSR-CR	0.01262	0.00601
SSR-DE	0.00801	0.00575

Notes. Standard errors are in parentheses. SSR-CR: sum of squared residuals for creation series. SSR-DE: sum of squared residuals for destruction series. P.A.M.: partial-adjustment model (σ_I not identified in this case). Quadratic: quadratic hazard model.

shows that, once more, the increasing-hazard model outperforms the partial-adjustment model. Figure V shows the residuals of both models for creation and destruction: as in the case of net flows, the improvement of the quadratic hazard model is largely gained during sharp contractions and brisk expansions.

Column 2 in Table II also shows that the hazard function is again asymmetric and clearly increasing. This hazard function does not differ significantly from that estimated using net flows (see Table I).

V. EXTENSIONS

In this section we outline additional features of the adjustment hazards approach presented in the previous sections as well as extensions of it.

V.1. Microeconomic Data

In our application and example we only used aggregate data. However, one of the main virtues of the approach we propose is that it provides a structural framework to use microeconomic data for improving the characterization and forecast of aggregate variables. Equation (6) can be run directly if information about the path of the moments of the cross-section distribution is available.
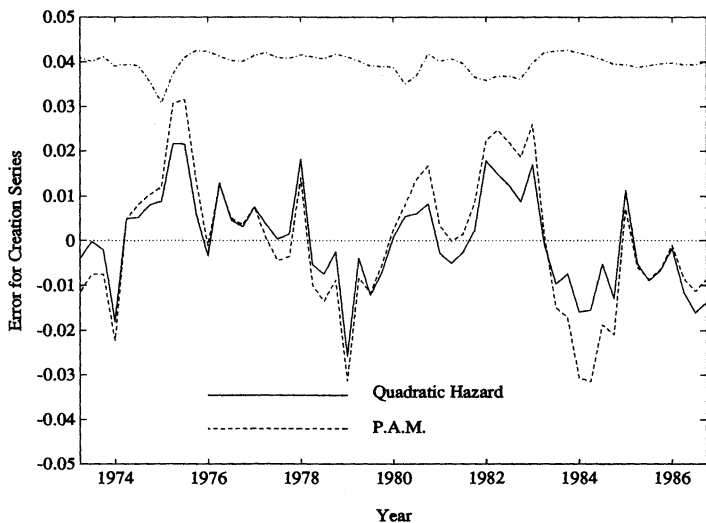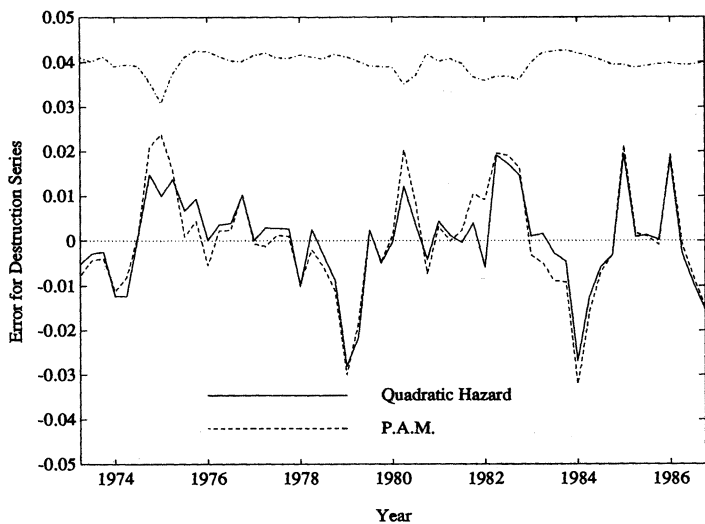
FIGURE Va

FIGURE Vb

Even though we do not have this information, we illustrate the procedure by integrating manufacturing aggregate and two-digit SIC data, as in Caballero and Engel [1992b]. We use BLS net flows and hours data for the period 1961:1–1983:1, and construct a proxy

for the path of the moments of the cross-sectional distribution under the assumption that they are proportional to the moments of the two-digit cross-sectional distribution.[24] We construct an estimating equation with the first five terms of equation (6) and run it by simple ordinary least squares. The dashed-dotted line in Figure III presents the estimated hazard function that, despite the approximations involved and the different period considered, is remarkably similar to those presented in the previous sections.

V.2. Heterogeneous Hazard Functions

For simplicity, we assumed that hazard functions are similar across units. This can be relaxed easily. If sectoral data are available, then the hazard function for every sector can be estimated separately. Alternatively, we can approximate the evolution of the aggregate by that of an economy with a hazard function equal to a weighted average of sectoral hazard functions.[25] These weights are determined not only by the fraction of firms with a given sectoral hazard function, but also by the corresponding ergodic distributions.

For example, consider the simplest case, where half the firms have a constant hazard equal to λ_L and the other half a constant hazard equal to λ_S, $\lambda_L > \lambda_S$. Then the "representative hazard function" $\Lambda(z)$, evaluated at zero will be near to λ_L, since firms with large adjustment hazards are more likely to be at $z = 0$. An analogous argument shows that, for large absolute deviations, $\Lambda(z)$ is near λ_S, since firms with small adjustment hazards are more likely to have large deviations. It follows that $\Lambda(z)$ is decreasing in the magnitude of the deviation, even though individual hazards are constant.

The main conclusion that follows from the preceding example is valid in general. Differences in bandwidths introduce a bias against the increasing hazard property. This provides an explanation for the asymmetry present in the hazard functions we estimated. If heterogeneity in firms' hazards is more prevalent in the hiring region, then we may expect there to be a range of deviations to the left of the return point for which the hazard function decreases with the size of the deviation.

---

24. The proportionality factor is obtained by choosing the value of v in the estimated hazard $\Lambda(vz)$ which minimizes the distance of this hazard from the hazard estimated using only aggregate data.

25. This approximation only captures the average across sectors; it is inappropriate for studying higher moment phenomena specific to any given sector.

V.3. Generalized Jump Functions

We assumed that firms set $z_{it} = 0$ every time they adjust, thereby reaching the employment level that is optimal if adjustment costs were momentarily removed. This can be relaxed to incorporate other constant as well as stochastic return points. The estimation strategy described in subsection IV.2 can be easily adapted to handle these cases if the return points are common to all agents jumping from the same position.

Generalizations of the jump function can also be used to extend our framework to other types of adjustment policies. Expressing the jump function as $-zJ(z)$, and using continuous time to simplify the notation, we can rewrite equation (1) as

$$(12) \qquad dE_t = -\int_{-\infty}^{\infty} zJ(z)\Lambda(z) f_t(z) \, dzdt.$$

This expression is quite general. For example, by letting $\Lambda(z)dt = 1$, it describes models of convex adjustment costs, with $J(z)$ representing the fraction of its deviation by which a unit at z adjusts.

Equation (12) shows that if we use micro data to obtain the moments of the cross-sectional density and then run a nonlinear regression (based on (12) or its discrete counterpart (6)), $J(z)$ and $\Lambda(z)$ cannot be identified separately: all we can estimate is their product. Yet if we consider estimation procedures—such as the one we use in this paper—that take into account the effect of shocks beyond the period where they occur, then $J(z)$ and $\Lambda(z)$ can be estimated separately, since different decompositions of $J(z)\Lambda(z)$ lead to different evolutions of the cross-sectional distribution and therefore different dynamics.[26]

VI. CONCLUSIONS

Typically, microeconomic units do not adjust continuously to all the shocks they perceive. Furthermore, the probability of adjusting is likely to depend on the unit's deviation from its target level. We characterize this behavior in terms of an adjustment hazard function, and study how aggregate and idiosyncratic shocks, filtered through this function and weighted by the cross-sectional density, determine the evolution of the aggregate.

Although the family of models we propose can be used to characterize the adjustment hazard without any commitment

---

26. However, this is not the case in the constant hazard model.

about its shape, we argue that realistic microeconomic descriptions should consider hazard functions that are eventually increasing with respect to the distance of the unit from its target. This has distinctive implications for aggregate relationships, which become nonlinear and exhibit complex dynamics. We estimate a hazard function for net and gross flows in U. S. manufacturing employment/jobs, and find evidence of the increasing-hazard property. More importantly, this model outperforms the partial-adjustment or quadratic-adjustment-cost–representative-agent model, especially during sharp recessions and brisk expansions.

One advantage of modeling the behavior of actual units and their cross section instead of using a representative-agent model, is that microeconomic information can be integrated into aggregate models naturally. Even though we have exploited this characteristic of adjustment-hazard models using only data that were far from fully disaggregated, this should play an important role in future applications of this methodology.

In sum, we provide a simple framework to understand the aggregate implications of a wide variety of realistic microeconomic policies. Our preliminary exploration of U. S. manufacturing employment data suggests that the nonlinearities and complex dynamics uncovered by this framework are highly relevant for applied work.

DEPARTMENT OF ECONOMICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY AND NBER
JOHN F. KENNEDY SCHOOL OF GOVERNMENT, HARVARD UNIVERSITY AND DEPARTAMENTO DE INGENIERIA INDUSTRIAL, UNIVERSIDAD DE CHILE

REFERENCES

Bar-Ilan, Avner, and Alan Blinder, "The Life Cycle Permanent Income Model and Consumer Durables," NBER Working Paper No. 2149, 1987.
Barro, Robert J., "A Theory of Monopolistic Price Adjustment," *Review of Economic Studies*, XXXIX (1972), 17–26.
Beaulieu, J. Joseph, "Aggregate Movements in Durable Goods with Fixed Costs of Adjustment," Unpublished manuscript, Massachusetts Institute of Technology, 1991.
Benabou, Roland, "Optimal Price Dynamics and Speculation with a Storable Good," *Econometrica*, LVII (1989), 41–81.
Bertola, Giuseppe, and Ricardo J. Caballero, "Kinked Adjustment Costs and Aggregate Dynamics," in *NBER Macroeconomics Annual 1990*, O. J. Blanchard and S. Fischer, eds. (Cambridge, MA: The MIT Press, 1990).
Bils, Mark, "The Cyclical Behavior of Marginal Cost and Price," *American Economic Review*, LXXVII (1987), 838–55.
Blinder, Alan S., "Retail Inventory Behavior and Business Fluctuations," *Brookings Papers on Economic Activity*, 2 (1981), 443–505.
Bresnahan, Timothy F., and Valerie A. Ramey, "The Nature of Output Fluctuations: Evidence from Plant Level Data," unpublished manuscript, 1991.
Caballero, Ricardo J., "Fuzzy Preferences and Aggregate Consumption," Columbia University mimeo, 1991.

——, "A Fallacy of Composition," *American Economic Review,* LXXXII (1992), 1279–92.

——, "Durable Goods: An Explanation for their Slow Adjustment," *Journal of Political Economy,* (1993), forthcoming.

Caballero, Ricardo J., and Eduardo M. R. A. Engel, "Dynamic (S,s) Economies," *Econometrica,* LIX (1991), 1659–86.

Caballero, Ricardo J., and Eduardo M. R. A. Engel, "Heterogeneity and Output Fluctuations in a Dynamic Menu Cost Economy," *Review of Economic Studies,* LX (1992a), 95–120.

Caballero, Ricardo J., and Eduardo M. R. A. Engel, "Beyond the Partial Adjustment Model," *American Economic Review, Papers and Proceedings,* LXXXII (1992b), 360–64.

Caballero, Ricardo J., and Eduardo M. R. A. Engel, "Price Rigidities, Asymmetries, and Output Fluctuations," NBER Working Paper No. 4091, 1992c.

Calvo, Guillermo A., "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics,* XII (1983), 383–98.

Caplin, Andrew S., "Variability of Aggregate Demand with (S,s) Inventory Policies," *Econometrica,* LIII (1985), 1395–1409.

Caplin, Andrew S., and John Leahy, "State-Dependent Pricing and the Dynamics of Money and Output," *Quarterly Journal of Economics,* CVI (1991), 683–708.

Caplin, Andrew S., and Daniel Spulber, "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics,* CII (1987), 703–26.

Davis, Steven J., and John Haltiwanger, "Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications," in *NBER Macroeconomics Annual 1987,* O. J. Blanchard and S. Fischer, eds. (Cambridge, MA: The MIT Press, 1990).

Davis, Steven J., and John Haltiwanger, "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics,* CVII (1992), 819–64.

Eberly, Janice C., "Adjustments of Consumers' Durables Stocks: Evidence from Automobile Purchases," unpublished manuscript, 1992.

Engel, Eduardo M. R. A., "Three Essays on Dynamic (S,s) Economies," Ph.D. thesis, Massachusetts Institute of Technology, 1991.

Grossman, Sanford J., and Guy Laroque, "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica,* LVIII (1990), 25–51.

Hamermesh, Daniel S., "Labor Demand and the Structure of Adjustment Costs," *American Economic Review,* LXXIX (1989), 674–89.

Rotemberg, Julio J., "The New Keynésian Microfoundations," in *NBER Macroeconomics Annual 1987,* O. J. Blanchard and S. Fischer, eds. (Cambridge, MA: The MIT Press, 1987).

Sargent, Thomas J., "Estimation of Dynamic Labor Demand Schedules under Rational Expectations," *Journal of Political Economy,* LXXXVI (1978), 1009–44.

Shapiro, Matthew D., "The Dynamic Demand for Capital and Labor," *Quarterly Journal of Economics,* CI (1986), 513–42.

Sheshinski, Eytan, and Yoram Weiss, "Inflation and Costs of Price Adjustment," *Review of Economic Studies,* XLIV (1977), 287–303.