Job Turnover and Policy Evaluation: A General Equilibrium Analysis

Hugo Hopenhayn

Northwestern University and Universitat Pompeu Fabra

Richard Rogerson

University of Minnesota

Recent empirical work indicates that job creation and destruction rates are large, implying significant amounts of job reallocation across firms. This paper builds a general equilibrium model of this reallocation process, calibrates it using data on firm-level dynamics, and evaluates the aggregate implications of policies that interfere with this process. We find that a tax on job destruction at the firm level has a sizable negative impact on total employment: a tax equal to 1 year's wages reduces employment by roughly 2.5 percent. More striking, however, are the welfare consequences: the cost in terms of consumption of this same tax is greater than 2 percent. The mechanism through which this welfare loss arises is apparently a decrease in average productivity, since this policy results in a decrease in average productivity of over 2 percent.

We have benefited from the comments of seminar participants at Arizona, Board of Governors, Buffalo, Columbia, Cornell, Maryland, the Minneapolis Fed, the NBER Summer Institute, Northwestern, Rice, Rochester, Santa Cruz, Stanford, Texas A&M, Toronto, Western Ontario, and Yale and from conference participants at the 1990 annual meetings of the Society for Economic Dynamics and Control in Minneapolis; the conference on Macroeconomics and Labor Markets held in Madrid, August 1990; the conference on Labor Markets and Wage Formation held in Amsterdam, January 1992; and the NBER Economic Fluctuations Conference in Palo Alto, Calif., February 1992. In particular, we would like to thank two anonymous referees, Jeremy Greenwood, and Ariel Pakes for many useful comments; Steve Davis and John Haltiwanger for providing us with data; and Les Reinhorn for research assistance. Support from the National Science Foundation is gratefully acknowledged.

[*Journal of Political Economy*, 1993, vol. 101, no. 5]
 0022-3808/93/0105-0004$01.50

I. Introduction

The goal of this paper is to use recent advances in the theoretical and empirical study of firm-level dynamics to examine the qualitative and quantitative impact of government policies that make it costly for firms to adjust their employment levels. Recent empirical work has analyzed several large data sets that track the evolution of individual firms over time.[1] Although details of the data sets and findings do differ across studies, one stylized fact that emerges is that the volume of job creation and destruction at the level of the individual firm is very large and that the vast majority of it is not accounted for by aggregate variables.

This finding is of particular interest given that cross-country comparisons indicate substantial differences in the relevant regulatory environments in which firms make employment decisions. Examples of regulations include legislated severance payments, plant closing legislation, advance notice, and various indirect bureaucratic costs (see Lazear [1990] for some cross-country comparisons). Given the amount of job creation and destruction taking place, one is naturally led to inquire about the effects of these types of regulations. In fact, there has recently been much discussion focusing on the role of labor market regulation as an explanatory factor for the differential performance of American and European labor markets over the last 20 years.

Recent theoretical work has produced several equilibrium models of industries that stress heterogeneity at the firm level.[2] In this paper we extend the industry equilibrium model of Hopenhayn (1992) to a general equilibrium setting. In a stationary equilibrium of this model, the aggregate properties of the economy are constant over time, although individual firms are continually adjusting, by expanding, contracting, starting up, or closing down. This environment stresses the heterogeneous development of firms and provides a natural setting in which to analyze policies that affect firm-level adjustments. We illustrate its usefulness for analyzing the types of policies mentioned above by considering the consequences of a policy that taxes firms for job destruction.

It is important to stress the role of recent empirical work in this exercise. Policy outcomes are strongly affected by the environment in which firms are operating. For example, in the case of dismissal

[1] Important contributions to this literature include Birch (1981, 1987), Dunne, Roberts, and Samuelson (1986, 1987, 1988, 1989), Leonard (1987), Evans (1987*a*, 1987*b*), Pakes and Ericson (1987), Davis and Haltiwanger (1988, 1990), and Troske (1989).

[2] Important contributions to this literature include Jovanovic (1982), Lambson (1988), Ericson and Pakes (1989), and Hopenhayn (1992).

costs, intuition suggests two opposing effects. On the one hand, firms may be less likely to dismiss workers in response to adverse shocks, possibly waiting for their situation to improve. On the other hand, given that dismissals are costly, firms may be less likely to hire workers in response to positive shocks, possibly waiting to see whether the situation persists before committing to hiring additional workers. The overall effect on aggregate employment seems ambiguous, depending on the stochastic structure of firm-level shocks. This being the case, evidence on the firm-level stochastic environment is necessary.

This work is closely related to a recent paper by Bentolila and Bertola (1990), who analyze the labor demand of a single monopolist subject to hiring and firing costs. Three differences are that their model abstracts from the entry/exit phenomenon, their calibration uses aggregate rather than firm-level data, and their analysis is partial equilibrium. Our results for the effects of a job destruction tax on employment are quite different from those obtained by Bentolila and Bertola: whereas they found that a dismissal cost actually increases long-run employment, we find that a tax on dismissals equal to 1 year's wages reduces long-run employment by roughly 2.5 percent. Furthermore, because we analyze the problem in a general equilibrium framework, the welfare consequences of these policies can also be analyzed. Our findings are quite striking: a tax equal to 1 year's wages reduces utility by over 2 percent measured in terms of consumption. By way of contrast, we note that it is common to find welfare costs on the order of a small fraction of a percent in most contexts.

An important channel through which this welfare reduction operates is average labor productivity, which is reduced by more than 2 percent by the introduction of this policy. Although the employment effects are of definite interest, we believe that the finding of such large welfare costs stands as one of the most important findings of our study. Policies that interfere with the job creation/destruction process are apparently quite costly.

This paper is intended as a first step in addressing the issues outlined above, and it is important to note several qualifications. First (as is also the case in Bentolila and Bertola), the effect of such policies on the nature of implicit or explicit labor contracts is ignored. Second, the analysis carried out here addresses only the long-run or steady-state effects of these policies and does not consider the short-run response of an economy to changes in policies. Third, for computational reasons, the current analysis excludes physical capital. Intuition suggests that if it were included the employment effects that result would be larger, since the types of policies studied presumably create an incentive for firms to substitute capital for labor. Finally, we focus here on only the costs associated with certain policies; we do not attempt to measure benefits that may be associated with such policies.

II. Model

The framework described below is designed for the purpose of studying a competitive economy that is in a stationary or long-run equilibrium. In this equilibrium, individual firms will be undergoing change over time, with some of them expanding, others contracting, some closing down, and others starting up. Despite all this change at the level of the individual firm, however, all aggregate variables—such as price, employment, output, and the number of firms—will be constant over time.

Each firm has a stochastic production function using labor as its only input. If a firm employs n_t workers in period t, when the output price is p_t it receives period t profits equal to

$$p_t f(n_t, s_t) - n_t - p_t c_f - g(n_t, n_{t-1}).$$

Several elements require some elaboration. The wage rate has been normalized to one and hence does not appear explicitly. The variable s_t is a firm-specific shock to production opportunities in period t. This shock takes values in R_+ and follows a first-order Markov process described by a function $F(s, s')$, with the interpretation that for each current value of the shock, denoted by s, $F(s, \cdot)$ is the distribution function for next period's value of the shock, denoted by s'. The shocks are independent across firms, but each firm's shock will evolve according to the same function F. The term c_f is a fixed operating cost (denominated in units of output) incurred by the firm in each period in which it remains in the market. As will become clearer later, the role of this fixed cost is to make it meaningful to talk about a firm exiting the market, as distinct from temporarily producing zero output. The function g captures the presence of adjustment costs and is included in the specification because the policy experiments to be studied later can be represented as changes in the g function. For example, the case in which a firm must make a fixed payment τ for every job that it destroys implies that g will take on the form

$$g(n_t, n_{t-1}) = \tau \cdot \max(0, n_{t-1} - n_t).$$

We now consider the decisions made by each firm and in particular the timing of these decisions. Figure 1 shows the sequence of decisions made by both incumbent and new firms. We begin with the period t decisions made by a firm that was in operation in period $t - 1$, at which time the firm had a shock equal to s_{t-1} and employed

Incumbent begins period t with (s_{t-1}, n_{t-1})

Exit Decision

Exit	Stay
receive $-g(0, n_{t-1})$ this period	find out value of s_t
zero in all future periods	make employment decision n_t
	receive $p_t f(n_t, s_t) - n_t - g(n_t, n_{t-1}) - p_t c_f$
	↓
	repeat next period

FIG. 1.—Timing of decisions

n_{t-1} workers. At the start of period t, prior to receiving any new information, the firm must first make a decision about whether to remain in the productive sector. If the firm exits, it implicitly chooses current employment equal to zero and must pay the adjustment costs $g(0, n_{t-1})$ associated with this choice, but it avoids paying the fixed operating cost c_f.[3] If a firm exits, it disappears from the model, receiving profits of zero in all future periods. If a firm chooses not to exit, then it incurs the fixed cost c_f and observes the current value of its shock, s_t. The firm then chooses current employment, produces output, and sells it at the period t price. This process is repeated next period.

Because of the nature of the exit decision, firms whose prospects look sufficiently poor (because of their last-period realization of s and the serial correlation in this process) exit to avoid the fixed cost. If there were no fixed cost, then the firm would not have to exit and could simply wait for the possibility of better times (higher realization of s), even if that implies output of zero in the immediate future.

Next consider the decisions made by potential entrants in period t. We assume that there is a large number (in fact a continuum) of ex ante identical potential entrants in each period. Entrants incur a one-time cost of c_e, again denominated in units of output. Once this cost has been paid, the entrant is in the same position as an incumbent that has chosen to remain in the productive sector and had zero employees in the previous period. Each new entrant receives its current value of s as a draw from the distribution ν. These draws are independently and identically distributed across entering firms, and

[3] This outcome explicitly assumes that if a firm closes down it still has to make good on all its obligations. It is of some interest to consider relaxing this feature.

the distribution ν is the same in each period and is independent of the number of entering firms. After the initial period, entering firms evolve in the same fashion as incumbent firms. Finally, all firms behave so as to maximize the expected discounted present value of profits, net of entry costs.

Having described the technology available to the economy, we still need to describe preferences and endowments. There is a continuum of identical agents, uniformly distributed over the unit interval, with preferences defined by

$$\sum_{t=1}^{\infty} \beta^t [u(c_t) - v(n_t)],$$

where c_t and n_t are consumption and labor supply in period t, respectively. Consumption is restricted to be nonnegative, and labor supply is restricted to be either zero or one. This last feature is included so that the number of employees at a firm is well defined. Following Hansen (1985) and Rogerson (1988), we assume that individuals choose employment lotteries and have access to markets to diversify idiosyncratic risk. This implies that the economy behaves as though there were a representative agent with preferences defined by

$$\sum_{t=1}^{\infty} \beta^t [u(c_t) - aN_t],$$

where N_t is the fraction of individuals who are employed in period t. The ownership of the technology is assumed to be uniformly distributed across the population, and profits are shared equally in equilibrium. Also, in the policy analysis conducted later in the paper, the revenues raised from taxing firms for job destruction are redistributed uniformly to all individuals as a lump-sum payment.

This completes the specification of the model. This model attributes all firm-level uncertainty to firm-level supply shocks. There is, however, an interpretation of this same structure in which the disturbances reflect shocks on the demand side. In this alternative structure, firms produce differentiated products, and the distribution of consumer tastes across differentiated products is stochastic over time. Although a firm's ability to physically produce output is constant over time, the value of this product is not. The production function is interpreted as specifying output in efficiency units that reflect the distribution of tastes in the market, and price corresponds to the price of an efficiency unit of output rather than the price of a physical unit of output.

III. Equilibrium

A. *Notation*

We begin by examining the decision problem of a firm in more detail. In anticipation of a stationary equilibrium, a constant output price of p is assumed. Specifically, consider a firm that employed n workers last period, decided to remain in the industry for the current period, and has received a new value for its shock equal to s. The Bellman equation corresponding to the firm's decision problem at this point is

$$W(s, n; p) = \max_{n' \geq 0} \{pf(n', s) - n' - pc_f - g(n', n) + \beta \max[E_s W(s', n'; p), -g(0, n')]\},$$

where E_s denotes expectations conditional on the current value of s, and s' denotes next period's (random) value of s. For reasons that will soon become clear, it is convenient to list the stationary price level p as a parameter in the value function. This equation is entirely straightforward with the possible exception of the maximization operator that is nested on the right-hand side. This reflects the fact that the firm will make a decision about exiting at the beginning of the next period. Moreover, because there will be no additional information revealed between the current decision point and the time of the exit decision, the firm can determine now whether it will choose to exit at that time. Of course, the decision to exit at the beginning of next period is not independent of this period's employment decision. Conditional on this period's employment choice, the firm must evaluate the expected value of remaining in the productive sector, given by $E_s W(s', n'; p)$, and compare this with the present discounted value of profits associated with exiting, given by $-g(0, n')$.

Note that if the value function W is known, the value of entering gross of entry costs, W^e, can be computed by

$$W^e(p) = \int W(s, 0; p) \, d\nu(s).$$

The firm's decision problem produces two decision rules: one for the optimal choice of current employment and the other for the optimal stay/exit decision at the beginning of next period. We write them as $N(s, n, p)$ and $X(s, n, p)$, respectively, with the convention that $X = 1$ corresponds to exit and $X = 0$ corresponds to stay.[4]

[4] The exit decision introduces a nonconvexity into the firm's decision problem that precludes standard results on uniqueness of decision rules from being applied here. Nonetheless, it can be shown that the decision rules are generically unique. Details are available from the authors on request.

The state of an individual firm is fully described by the pair (s, n). The state of the economy is described by the distribution of the state variables for all individual firms. It is natural to express this as a measure over pairs (s, n), which we denote as $\mu(s, n)$. In the computations performed later in this paper, both s and n will be restricted to take on a finite number of values, in which case μ can be represented as a matrix, with the i-j element giving the total number (or mass) of firms that have their individual state variable equal to the pair (s_i, n_j).

The information introduced thus far is sufficient to trace the evolution of the economy over time, assuming that price is constant. In period t, at the point at which incumbents have made their stay/exit decision and new realizations of s have been drawn, let the incumbents be summarized by a measure μ, and let the mass of firms that enter be equal to M. Firms make optimal employment decisions using the decision rule $N(s, n; p)$, and at the beginning of next period some of them exit according to the decision rule $X(s, n; p)$. The aggregate state for period $t + 1$'s incumbents after exit decisions have been made and new information has been revealed will be given by some measure μ'. The transition from μ to μ' will be written as $\mu' = T(\mu, M; p)$. Note that the (constant) price p enters because it determines the decision rules used by firms. The operator T is linearly homogeneous in μ and M jointly, a property that is used later in the paper. In particular, note that T is not linearly homogeneous in μ alone; even if the economy begins with twice as many firms of each type, the economy does not end up with twice as many firms of each type one period later unless entry is also doubled.

The amount of output, Y, produced in a given period as a function of the variables μ, M, and p can be determined as

$$Y(\mu, M; p) = \int [f(N(s, n; p), s) - c_f] d\mu(s, n) + M \int f(N(s, 0; p), s) d\nu(s).$$

In the first integral, output (net of fixed operating costs) for a firm with state variable (s, n) is computed using the optimal employment rule N and then integrated over the distribution of incumbents. The second integral does the same for new entrants, the only difference being that all new entrants have a value of zero for last period's employment, and their distribution of idiosyncratic shocks is given by ν.[5]

[5] Note that we are assuming that a new entrant bears only the fixed cost of entry, c_e, and does not pay the cost c_f.

Several other aggregates can be defined similarly. An individual firm with last-period employment n and current shock s has expected adjustment costs for the next period given by

$$r(s,n;p) = [1 - X(s,n;p)] \int g(N(s',n';p),n')\,dF(s,s') + X(s,n;p)\cdot g(0,n'),$$

where $n' = N(s, n; p)$. Integration yields aggregate adjustment costs specified by $R(\mu, M, p)$. Labor demand and profits are given by

$$L^d(\mu,M,p) = \int N(s,n;p)\,d\mu(s,n) + M\int N(s,0;p)\,d\nu(s),$$

$$\Pi(\mu,M,p) = pY(\mu,M,p) - L^d(\mu,M,p) - R(\mu,M,p) - Mpc_e.$$

It is straightforward to show that Y, R, L^d, and Π are linearly homogeneous in μ and M jointly. Some notation is also necessary to describe the consumer problem. In a stationary state with constant prices and the interest rate satisfying $1/(1 + r) = \beta$, the individual optimization problem reduces to a static optimization problem of the form

$$\max u(c) - aN \quad \text{subject to}\ pc \leq N + \Pi + R,$$

where Π is profits and R is tax receipts. The solution to this problem can be written as $N = L^s(p, \Pi + R)$.

B. *Definition of Equilibrium*

A stationary equilibrium for the model introduced in Section II is given by the following definition.

DEFINITION. A stationary equilibrium consists of an output price $p^* \geq 0$, a mass of entrants $M^* \geq 0$, and a measure of incumbents μ^*, such that (i) $L^d(\mu^*, M^*, p^*) = L^s(p^*, \Pi(\mu^*, M^*, p^*) + R(\mu^*, M^*, p^*))$, (ii) $T(\mu^*, M^*; p^*) = \mu^*$, and (iii) $W^e(p^*) \leq p^* c_e$ with equality if $M^* > 0$.

These conditions require little explanation. Condition i states that demand must equal supply in the labor market. Condition ii states that the state of the economy must be such that the optimal actions of firms cause this state to be reproduced in each period. Condition iii states that entering firms must be willing to enter: if M^* is strictly positive, $W^e(p^*)$ must be equal to $p^* c_e$. Note that in equilibrium $W^e(p^*)$ cannot be strictly bigger than $p^* c_e$ because of the assumption that there is an unlimited supply of potential entrants. It is possible, however, for an equilibrium to entail $W^e(p^*) < p^* c_e$, although in this case there will not be any entry, and consequently, by condition ii, there must also be no exit.

C. *Existence and an Algorithm for Finding an Equilibrium with Entry and Exit*

This section outlines an algorithm for finding an equilibrium with entry and exit. There are two reasons for doing so. First, this algorithm is the one used later in the paper to numerically compute stationary equilibria for the model. Second, the algorithm provides insights into the workings of the model and, in particular, highlights some of the interactions between features of the model specification and computational complexity.

The following assumptions will be made.

ASSUMPTION 1. The production function $f(n, s)$ is continuously differentiable and strictly concave in n for each value of s. The function $f_1(n, s)$ is strictly positive and increasing in s and satisfies $\lim_{n\to 0} f_1(n, s) = \infty$.

ASSUMPTION 2. The income effect on labor supply is negative; that is, L_2^s is negative.

ASSUMPTION 3. The function F is continuous and decreasing in its first argument, and ν has a continuous cumulative distribution function.

As was noted following the definition of a stationary equilibrium, there are two different forms that the equilibrium may take: with entry and exit or without. If an equilibrium with entry and exit exists, then it is the only such equilibrium, and there are no equilibria without exit and entry. On the other hand, if an equilibrium without entry and exit exists, then generically there will be a continuum of equilibria. As a practical matter, the case with entry and exit is of greater interest here, since the data used to calibrate the model in a later section indicate that significant amounts of entry and exit are taking place. As a result, the discussion that follows focuses on the case in which entry and exit occur in equilibrium. For some discussion of the case without entry and exit, see Hopenhayn (1992).

The algorithm consists of three steps. The first step uses condition iii of equilibrium to find the price p^*, the second step uses condition ii of equilibrium to find μ^* up to a scale factor, and the third step uses condition i to determine the scale factor, which turns out to equal M^*.

The first step of the algorithm finds the unique value of p that is consistent with entry in equilibrium. For any given p, one can compute the value function $W(s, n; p)$ and hence $W^e(p)$. The function W

is strictly increasing and continuous in p, and hence so is W^e. It follows that there is exactly one positive value of p that can satisfy $W^e(p) = p \cdot c_e$, and this must be the value of p if a stationary equilibrium with entry and exit exists.[6] Call this value p^*.

The second step of the algorithm determines whether or not an equilibrium with entry and exit exists and, if so, finds μ^* up to a scale factor. The decision rules $N(s, n; p^*)$ and $X(s, n; p^*)$ can be used to compute the transition function T. A stationary equilibrium requires a pair (μ^*, M^*) such that μ^* is a fixed point of $T(\mu, M^*; p^*)$. Given p^* and M^*, this operator is affine and has at most one fixed point. Moreover, the linear homogeneity of T in μ and M implies that if μ^* is a fixed point when entry is one, then $M\mu^*$ is a fixed point when entry equals M, for all $M > 0$.

Under what conditions will this operator have a fixed point? In a stationary state the inflow of firms must equal the outflow of firms. If there is a mass of entry equal to M in each period, then there must be a mass of exit of M in each period. Furthermore, since there are new firms entering each period, all firms must eventually exit, or else the number of firms would be growing over time. It turns out, however, that it is not sufficient to simply have all firms exit eventually; if the exit occurs too far into the future, then the constant inflow will still cause the industry to grow over time. The condition that does ensure the existence of a fixed point is that the expected age at exit be finite for any entering firm. In practice, this amounts to a joint condition on the decision rules and the exogenous stochastic process. These conditions will not be developed here; the interested reader can refer to Hopenhayn (1992) for a discussion in the context of a simpler model.

Consider a unit mass of entry at some time, and let these firms evolve according to the decision rules $N(s, n; p^*)$ and $X(s, n; p^*)$. Let λ_t be the fraction of these firms that are still around t periods later, and let ϕ_t be the distribution of the state variables of these firms after the exit decision has been made.[7] If there is a unit mass of entry in each period, then with price equal to p^* the state of the economy must converge to $\Sigma_{t=1}^{\infty} \lambda_t \phi_t$.

The condition that this sum be finite is exactly the condition that expected age at exit be finite. If the sum is finite, it is the unique

[6] It is easy to show that $p = 0$ is also always a solution to this equation. To see that there is a unique positive solution, note that expected discounted profits are linearly homogeneous in the wage and price, and decreasing in the wage. Assume that p and λp are both solutions ($\lambda > 1$). Then expected discounted profits with price λp and wages equal to λ are equal to $\lambda p c_e$, and when wages are reduced to one, expected discounted profits will exceed $\lambda p c_e$.

[7] Note that ϕ_t will satisfy $\phi_t = T(\phi_{t-1}, 0, p^*)$.

fixed point of $T(\mu, 1; p^*)$, and we denote it by $\hat{\mu}$. It is immediate from the linear homogeneity of T that if entry were M in each period rather than one, then the sequence would converge to $M\hat{\mu}$. Hence, this procedure produces a continuum of pairs (μ, M), parameterized as $(M\hat{\mu}, M)$ for $M > 0$, that replicate the distribution μ.

The third step determines the scale factor M^*. Assume that a fixed point $\hat{\mu}$ has been found. The equilibrium must also satisfy condition i, and hence M must be chosen to satisfy

$$L^d(M\hat{\mu}, M; p^*) = L^s(p^*, M(\hat{\Pi} + \hat{R})),$$

where $\hat{\Pi} = \Pi(\hat{\mu}, 1, p^*)$ and $\hat{R} = R(\hat{\mu}, 1, p^*)$. As noted previously, the left-hand side of this expression is linearly homogeneous in M, and by assumption 2 the right-hand side is strictly decreasing in M, implying a unique value of M that satisfies this equation.[8] Under the assumption that $\hat{\mu}$ exists, this procedure determines the unique stationary equilibrium.

It is possible, however, that $\hat{\mu}$ does not exist. In this case there is no equilibrium with entry and exit, and typically there will exist a continuum of equilibria, indexed by the size of the market. The empirical work to follow involves only the case in which entry and exit do occur in equilibrium, so we do not discuss further the details of the other case.

As a final remark in this section, we point out how one change in the specification of the model affects the algorithm above. It has been assumed that the distribution of new entrants, ν, is independent of the amount of entry or, similarly, that the cost of entry is independent of the amount of entry. Allowing for the possibility that the quality of entrants is decreasing in the amount of entry will destroy the linear homogeneity of the T operator and hence require a different (and more complicated) algorithm from the one used above. Nonetheless, incorporating these types of changes is feasible, and Hopenhayn (1990) does consider such cases in a slightly different context.

IV. Benchmark Model

A. *Specification*

The two preceding sections have laid out a model for which a stationary equilibrium has the property that aggregate variables are constant over time even though at the level of the individual firm there is considerable change over time. At any point in time there are some

[8] This argument assumes that $\hat{\Pi} + \hat{R}$ is strictly positive in equilibrium. This is necessarily true when adjustment costs are zero, but even if they were negative the expression above would still have at least one solution.

firms expanding, some contracting, some exiting, and others entering. Moreover, there is something "good" about this process: efficient firms are expanding and inefficient firms are contracting, possibly exiting, and being replaced by firms that are more productive in an expected sense. This provides the setting in which we carry out the policy experiments outlined in the Introduction.

To carry out a quantitative analysis, we must specify the model more explicitly, choosing functional forms and assigning parameter values. This section specifies the benchmark model used for the policy experiments and reports some of its properties. The next section discusses how the model is calibrated.

The benchmark model consists of the following functional forms:

$$f(n, s) = sn^{\theta}, \quad 0 \leq \theta \leq 1,$$

$$g(n_t, n_{t-1}) = 0,$$

$$\log(s_t) = a + \rho \log(s_{t-1}) + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma_\epsilon^2), a \geq 0, 0 \leq \rho < 1,$$

$$u(c) = \ln(c), \quad v(n) = An, \quad A > 0.$$

Note that these functional forms satisfy assumptions 1–3. A few comments are in order with respect to this choice of a benchmark case. The production function seems quite natural, so there is little to say about this choice. The assumption of no adjustment costs in the benchmark case is primarily a choice of convenience. As will be seen shortly, this specification simplifies the procedure for calibrating the model, although the procedure can accommodate other choices of g.

The choice of a process for s is more difficult. On the positive side, the specified process has the advantage of presenting a parsimonious representation, with the parameters corresponding to objects that are of intuitive interest given the nature of the policy experiments performed. For example, the parameter ρ is a measure of persistence in the s process, and it is expected that changes in the persistence of the shocks will have an impact on how much a firm is affected by legislated dismissal costs. If persistence is very high, then, loosely speaking, a firm expects that jobs created today will be around for a long time, and hence the dismissal costs will be strongly discounted and not play much of a role in the firm's decisions. Conversely, if shocks are not very persistent, then a firm will want to take the dismissal costs into account because there is a strong possibility that they will be incurred relatively soon.

The assumed process for s exhibits mean reversion, a property that is inherited by the endogenous stochastic process for firm size. This property is consistent with recent findings from firm-level data sets.

On the negative side, however, this process for s abstracts from some features that have been found in firm-level data sets. In particular, this process assumes that the level of persistence is independent of the current value of s, as is the variance of the innovation, ϵ_t. Studies have found that larger firms tend to display both more persistence and less variability than smaller firms do. Also, age effects have been noted by several studies. We note that it is feasible to include many of these features into the specification at little or no increase in computational complexity. It is our view, however, that as a first step it is preferable to start with this process for s as a way to economize on parameters that need to be chosen. The specification of preferences is the same as that employed by Hansen (1985) in his study of business cycles and has been used as a benchmark in many subsequent aggregate studies. In the neoclassical one-sector growth model, this specification implies that employment is constant along a balanced growth path.

With the functional forms discussed above, the firm's dynamic programming problem is considerably simpler since the state variable for last period's employment is no longer relevant. In fact, it is easy to show that the optimal decision rules for individual firm behavior imply

$$\log n_t = \frac{1}{1-\theta}(\log\theta + \log p + \log s_t),$$

$$X(s_t, n_t, p) = 1 \quad \text{if } s_t \leq s^* \text{ for some } s^*,$$

where, as indicated previously, X equal to one is interpreted as exit.

The exit rule instructs the firm to exit if s_t is below some reservation value, denoted by s^*. Given the process for s, this result should not be surprising: higher values of s indicate higher expected future values of s, and hence $E_s W(n', s')$ is increasing in s (and in the benchmark independent of n'). Although a closed-form solution for the employment decision rule is available, a closed-form solution for s^* is not. In combination, these two decision rules imply the following law of motion for employment of a surviving firm:

$$\log(n_{t+1}) = \frac{1-\rho}{1-\theta}\left(\log\theta + \log p + \frac{a}{1-\rho}\right) + \rho\log(n_{t-1}) + \left(\frac{1}{1-\theta}\right)\epsilon_t. \tag{1}$$

Note that the form of this relationship is effectively independent of the actual equilibrium price level. Estimates of this relationship using cross-sectional data on surviving firms will provide information on the parameters of the process for exogenous shocks, a point we return to in the next section. For future reference, two comparative static results about the effect of the two fixed costs on the stationary equilibrium for the benchmark model are noted. These results are proved in a more general setting in Hopenhayn (1992), so the reader is referred there for details.

PROPOSITION 1. If the stationary equilibrium involves exit and entry, then an increase in c_f will increase the price, the average size of a firm, and the exit rate.

This result is fairly intuitive. At the initial equilibrium price, if c_f increases, the value of s^* will increase, causing the exit rate to increase and the average size of a firm to increase. At the initial price, firms will no longer be willing to enter, causing the price to increase but only partially offsetting the initial effects. The next proposition gives a result for changes in the cost of entry.

PROPOSITION 2. If equilibrium involves entry and exit, then an increase in c_e will result in a higher price and a lower exit rate.

This result is also intuitive. At the initial price, if c_e increases, then exit drops to zero, causing the price to increase. This lowers the value of s^*, decreasing the exit rate.

B. *Calibration Procedure*

There are many procedures that could be specified to choose parameter values for the model specified above; we describe one method. Many of the parameter values are dependent on the length of a time period, which is set to 5 years. There are several reasons for this choice. First, the time interval of the *Census of Manufactures* data set, which we use here, is 5 years. Second, there is a desire to abstract from higher-frequency movements such as cyclical fluctuations, since the model abstracts from such features. Also, for the policy experiments carried out, it is desirable to abstract from temporary layoffs. Third, for computational reasons it is preferable to use a period of at least this length to ease the computational burden.

The procedure used to assign parameter values is in the same spirit as the one pioneered by Kydland and Prescott (1982) in the real business cycle literature. Loosely speaking, this amounts to using the same number of statistics as there are parameters to be assigned and choosing the parameters so that the model's equilibrium exactly matches the chosen statistics. Christiano and Eichenbaum (1988) have shown how this procedure can be formally cast as using method of moments estimators in an exactly identified system.

The parameter θ is equal to labor's share of total revenue, and the parameter β maps monotonically into the real interest rate, providing a method for assigning these two values.

Recall expression (1) for the law of motion for employment of surviving firms. Given a value for θ, in a regression of $\log(n_t)$ on a constant and $\log(n_{t-1})$ for surviving firms, the coefficient on $\log(n_{t-1})$ is an estimate of ρ, and the residual variance is equal to $\sigma_\epsilon^2/(1 - \theta)^2$.

The parameters that remain to be assigned are the two fixed costs, c_e and c_f, the constant a in the law of motion for the shock, and the initial distribution ν. In the absence of data that allow the price to be estimated, there is an identification problem with respect to the first three parameters. The basic issue is that the price of output and the idiosyncratic shock enter multiplicatively in the firm's objective function, and effectively one cannot disentangle the difference between a high price and a high average value of the idiosyncratic shock. Given the abstract nature of the model, it is not clear how one would use actual price data to determine an appropriate value for p^*, and hence the identification problem needs to be dealt with. This is accomplished by normalizing the stationary equilibrium price in the benchmark model to be unity and choosing the values of the model's remaining parameters to be consistent with this.

The statistics used to determine c_f and a are the cross-sectional average of log employment and the 5-year exit rate. Values for c_f and a cannot be determined analytically, but there is a unique choice of c_f and a that matches these two numbers. This follows from the fact that s^* is increasing in c_f and a and that, given s^*, the average size is increasing in a and independent of c_f.

Given values for all the other parameters, a choice of ν will determine the size distribution by cohort in the stationary equilibrium. Matching some part of this distribution provides a method to choose ν. We use the actual size distribution of firms aged 0–6 years and choose a value of ν so that our distribution of firms in their first and second periods matches this. We found that a uniform distribution on the lower part of the interval in which realizations of s lie produced a reasonable fit.

The value of c_e is chosen so that condition iii of the definition of equilibrium is satisfied with $p = 1$. It should be noted that as a practical matter, since entry and exit are nonzero in the actual data, the only cases that will be dealt with in the numerical exercises are those in which the stationary equilibrium involves entry and exit. Hence, the multiplicity of equilibria in the case without entry and exit is not an issue in the computations performed later in the paper. Finally, the value of A is chosen to produce an employment to population ratio equal to .6.

C. *Data and Parameter Assignments*

Having described the procedure for choosing parameter values, we still need to obtain the required measures from the empirical literature. The numerical work performed later in the paper will include sensitivity analysis to determine the effect of small changes in the values of the parameters.

It is relatively straightforward to choose values for those parameters that can be linked to more highly aggregated statistics. The discount rate is set to .8, which with a period interpreted as 5 years implies an annual real interest rate of roughly 4 percent. The parameter θ is set to .64, which is the value that has been used in most of the real business cycle literature.

For the remaining measures it is necessary to use firm-level data sets, of which there are several that provide all or some of the required information. These data sets all have their particular strengths and weaknesses, relating to such issues as the extent of coverage, information included, frequency of data, and the ability to detect mergers, acquisitions, and transfers. Parameter values for the benchmark case are based on data from the Longitudinal Research Data (LRD) file, covering the years 1972 and 1977.[9] Table 1 provides the relevant measures that are taken from these data.

D. *Equilibrium for the Benchmark Model*

With all the parameters determined, the benchmark model can be solved numerically. The numerical procedure used is relatively straightforward, and hence only a few details are provided here. Readers can contact the authors for more details and copies of the computer programs.

To compute the value functions, we discretize the state space and iterate on Bellman's equation. Grids of 250 points for employment and 20 points for the idiosyncratic shock were used. Log scales were used in both cases, and the shocks were picked so that the maximum employment level was equal to 5,000 workers. A discrete approximation was made to the first-order Markov process on $\log(s)$. Because s is exogenous, the choice of number of points in the grid for s is effectively an assumption. In the case of the employment variable this is not true: because n is endogenously chosen, one needs to be sure that the coarseness of the grid is not significantly affecting the results. Sensitivity analysis indicated that the choice of 250 points was adequate to guarantee this.

[9] The LRD is a national sample of manufacturing establishments, consisting of a sequence of contiguous 5-year panels beginning in the years 1963, 1967, 1972, 1977, and 1982. This data set has recently been utilized by Davis and Haltiwanger (1988, 1990). The initial year in each of these panels is used to produce the *Census of Manufactures*. The census data have been used by Dunne et al. (1986, 1987, 1988, 1989) in a series of investigations.

TABLE 1

A. Estimates Derived from the LRD

Serial correlation in log employment (5-year interval, survivors)	.93
Variance in growth rates (log difference, 5-year interval, survivors)	.53
Mean employment	61.7
Exit rate (5-year interval)	37%

B. Size Distribution for Firms Aged 0–6 Years

Employees	Share of Total Firms
1–19	.74
20–99	.18
100–499	.08
500+	.01

There are a number of statistics that can be reported in characterizing the equilibrium for this model, and some of them are contained in table 2.

Several properties emerge. First, note the statistics that are reported by size (part B). The size distribution of firms indicates that most firms are in fact quite small. However, the size distribution of employment indicates that although most firms are small, most employment is accounted for by larger firms. The mean firm size and the co-worker mean reported in part A of the table are supporting pieces of information. Although these statistics were not explicitly calibrated, they are in fact quite close to those reported in Birch (1987), Davis and Haltiwanger (1988), and Troske (1989). The size distribution of hiring and firing provides an expected pattern given that firm-level employment is following a mean-reverting process: most of the firing is done by larger firms and most of the hiring is done by smaller firms. Part A of the table shows that the average size of entering and exiting firms is quite small, which is also consistent with available evidence (see Dunne et al. 1986, 1987, 1988; Troske 1989).

The statistics related to cohorts indicate two patterns. First, the probability of exit is decreasing in age, and, second, the size distribution of firms is stochastically increasing in age; that is, the size distribution moves to the right as the age of the cohort increases. Both of these properties have been noted by empirical work in this area (see, e.g., Evans 1987*b*).

It is perhaps important to indicate how the reader should interpret these statistics. The model specification that has been chosen has relatively few parameters, and as mentioned earlier with respect to the process on firm-level employment, this specification will not match

TABLE 2

A. Summary Statistics for Benchmark Model

Average firm size	61.2
Co-worker mean	747
Variance of growth rates (survivors)	.55
Serial correlation in log n (survivors)	.92
Exit rate of firms	.39
Turnover rate of jobs	.30
Fraction of hiring by new firms	.15
Average size of new firm	7.5
Average size of existing firm	4.9

B. Size Distribution

	1–19	20–99	100–499	500+
Firms	.52	.37	.10	.01
Employment	.06	.24	.37	.33
Hiring	.05	.35	.41	.19
Firing	.12	.19	.34	.35
By cohort:				
1 period	.88	.12	.00	.00
2 periods	.54	.45	.01	.00
5 periods	.29	.58	.12	.01
10 periods	.20	.54	.20	.05
Hazard rates by cohort:				
1 period	.75			
2 periods	.32			
5 periods	.15			
10 periods	.10			

with the data equally well on all dimensions. Moreover, when the model is calibrated using a small set of empirical statistics, there are conceivably other dimensions along which the model may not fit particularly well. The statistics above are shown to indicate that the relatively simple structure used here with the calibrated parameter values does a reasonable job of matching several aspects of the relevant data, both qualitatively and quantitatively.

V. Results

This section reports the results from introducing an adjustment cost function of the form

$$g(n_t, n_{t-1}) = \tau \cdot \max\{0, n_{t-1} - n_t\},$$

with the interpretation that τ is a tax that the firm must pay for each job that is destroyed. The size of τ can be interpreted by comparison with other values in the model; with a period equal to 5 years and w

TABLE 3

Effect of Changes in τ (Benchmark Model)

	$\tau = 0$	$\tau = .1$	$\tau = .2$
Price	1.00	1.026	1.048
Consumption (output)	100	97.5	95.4
Average productivity	100	99.2	97.9
Total employment	100	98.3	97.5
Utility-adjusted consumption	100	98.7	97.2
Average firm size	61.2	61.8	65.1
Layoff costs/wage bill	0	.026	.044
Job turnover rate	.30	.26	.22
Serial correlation in log(n)	.92	.94	.94
Variance in growth rates	.55	.45	.39

normalized to one, a value of τ equal to .1 corresponds to 6 months' wages, and a value of τ equal to .2 corresponds to 1 year's wages. Table 3 reports how the equilibrium is affected when τ takes on the values of .1 and .2. To facilitate comparison with the benchmark model, some of the earlier values are repeated in this table, and for cases in which interest is primarily in relative changes, an index has been created in which the $\tau = 0$ values are set equal to 100. We do not report changes for all the variables in table 2; changes in the size distribution of firms and other distributional statistics were relatively minor.

Qualitatively, these results are quite intuitive. The tax on dismissals causes firms to be more cautious about job creation and thereby also reduces the need for job destruction, with the net result that firms end up making fewer adjustments to their labor forces. These effects show up in the new equilibria: as τ increases, the serial correlation in log employment increases, whereas the variance in growth rates and the job turnover rate decrease.

The results indicate a fairly strong trade-off between the average duration of a job and the total number of jobs. As one moves from $\tau = 0$ to $\tau = .2$, the job destruction rate decreases by 8 percent, whereas total employment goes down by roughly 2.5 percent. To the extent that $\tau = .2$ is a reasonable description of the magnitude of legislated severance payments in several countries (see Lazear [1990] for more country-specific details), we believe that the 2.5 percent decrease in employment is very significant. Although the trade-off between total employment and job duration is an interesting piece of information contained in table 3, the result concerning the efficiency costs of these policies is also striking. The figure for utility-adjusted consumption shows the amount by which consumption would have to be increased in order for utility to reach the same level attained when $\tau = 0$. This measure takes into account the fact that leisure is higher when τ is positive. When $\tau = .2$, consumption would have to increase by 2.8 percent in order to compensate individuals for the loss of utility, a magnitude that is quite impressive since in many contexts the utility consequences of distortions tend to be on the order of a fraction of a percent.

A closely related piece of information is the effect on productivity. As indicated in the table, productivity drops by 2.1 percent when τ is set to .2. This decrease accords well with intuition: the tax on job destruction creates a distortion that encourages firms to use resources less efficiently, with the result that productivity drops and fewer resources are devoted to the market sector of the economy. It is instructive to examine how the dismissal cost affects the decision rules of individual firms. In the benchmark model, current employment is determined entirely by the current value of the idiosyncratic shock. In the presence of a dismissal cost, current employment is also affected by last period's employment. With linear adjustment costs, it is well known that the decision rules have the property that, for $n_{t-1} \in [n_l(s_t), n_u(s_t)]$, current employment is equal to the last period's employment and is independent of n_{t-1} otherwise. The size of the band $[n_l(s_t), n_u(s_t)]$ provides some useful information about the extent to which resources are not allocated efficiently. Table 4 contains information about these bands for $\tau = .1$ and $\tau = .2$ for several values of the idiosyncratic shock. As expected, the bands are much wider for $\tau = .2$. Note also that the bands are quite large: when $\tau = .2$, the band width is typically more than one-third of the midpoint of the interval.

In the model with policy distortions, employment is always set such that the marginal product of labor (MPL) equals $1/p$. A natural way to document the misallocation of resources that arises when firms use the decision rules described above is to look at absolute deviations of MPL from $1/p$ in the stationary distribution. Table 5 provides information on this distribution. As is easily seen, deviations become significantly larger as τ increases. When $\tau = .2$, almost 90 percent of all firms have a deviation greater than 5 percent.

An additional point concerning the productivity numbers that is important to recognize concerns robustness. Whereas the magnitude of the employment effect is influenced by the form of preferences used, this is not true for the effect on productivity. In the algorithm used to find an equilibrium, preferences affect only the scale of activity; the stationary distribution of firms is determined up to a scale factor independently of preferences. Because the scale factor does not affect productivity, the productivity effects reported in the table are independent of preferences, which is also true for the change in

TABLE 4

Effect of τ on Decision Rules

	$\tau = .1$		$\tau = .2$	
log s	n_l	n_u	n_l	n_u
1.83	1.36	1.78	1.18	1.98
4.75	21.7	26.7	21.0	32.8
10.5	194	238	181	282
19.9	1,110	1,410	1,036	1,617
27.3	2,610	3,316	2,522	3,935

the job turnover rate. This independence results from the fact that the *aggregate* technology displays constant returns to scale. In view of the large welfare effects of these policies, we feel that it is misleading for attention to be focused exclusively on the employment consequences of dismissal costs. The message that emerges from this analysis is that it is very costly to distort the job creation/destruction process. Note also that the fraction of total payroll that is paid in dismissal costs is not particularly large: even when $\tau = .2$, they account for less than 5 percent of total payroll.

As noted above, the magnitude of the change in employment depends on preferences. As argued earlier, however, there is a good reason for the choice of preferences discussed above. Nonetheless, to illustrate the impact of alternative preferences, consider preferences with $u(c) = c^{\alpha}/\alpha$, so that $\log(c)$ corresponds to $\alpha = 0$. The change in employment associated with moving from $\tau = 0$ to $\tau = .2$ is a decrease of 3.4 percent when $\alpha = 1$. Using the fact that nonlabor income is .2034 and .2197, respectively, one can straightforwardly compute the change for other values as well.

It is obviously of interest to know how sensitive the results are to

TABLE 5

Absolute Deviations from MPL $= 1/p$

	Fraction of Firms within Interval	
Size of Deviation (%)	$\tau = .1$	$\tau = .2$
0–3	.30	.00
3–5	.45	.12
5–10	.15	.78
10–15	.00	.05
>15	.00	.05

small changes in some of the key parameters. We have investigated this by analyzing deviations of plus and minus 10 percent in the values of β, θ, ρ, σ_ϵ^2, and a. In each of these cases the model was not recalibrated; all other parameters were kept at their values in the benchmark model. These changes had very little impact on the effects of dismissal costs, and hence we do not report the results here.

VI. Conclusion

The objective of this paper has been to use recent theoretical and empirical work on firm-level dynamics to begin a quantitative evaluation of labor market policies that affect firm-level decisions about labor force adjustment. We found that a tax on job destruction significantly reduced steady-state employment. More important, we also found that these policies implied large welfare losses, resulting primarily from a significant decrease in average labor productivity. As was emphasized in the Introduction, there are a number of qualifications to be made concerning the interpretation of the specific results obtained. However, these qualifications can be viewed as providing a natural set of issues to be addressed by future work. The type of exercise carried out in this paper seems to provide some useful information about the quantitative significance of some of the model's features, thereby providing a useful link between theory and measurement in this area. A final point to note is that the analysis has focused entirely on the costs associated with particular policies and has not dealt with the benefits that might be associated with these policies. Obviously costs must be compared to benefits to assess the significance of their magnitude.

References

Bentolila, Samuel, and Bertola, Giuseppe. "Firing Costs and Labour Demand: How Bad Is Eurosclerosis?" *Rev. Econ. Studies* 57 (July 1990): 381–402.

Birch, David L. *The Job Generation Process.* Cambridge, Mass.: MIT Press, 1981.

———. *Job Creation in America: How Our Smallest Companies Put the Most People to Work.* Cambridge, Mass.: MIT Press, 1987.

Christiano, Lawrence, and Eichenbaum, Martin. "Is Theory Really Ahead of Measurement? Current Real Business Cycle Theories and Aggregate Labor Market Fluctuations." Working Paper no. 2700. Cambridge, Mass.: NBER, September 1988.

Davis, Steve, and Haltiwanger, John. "The Distribution of Employees and Establishments by Establishment Size: Patterns of Change and Comovement in the United States." Manuscript. Chicago: Univ. Chicago, 1988.

———. "Gross Job Creation and Destruction: Microeconomic Evidence and

Macroeconomic Implications." In *NBER Macroeconomics Annual*, vol. 5, edited by Olivier J. Blanchard and Stanley Fischer. Cambridge, Mass.: MIT Press, 1990.

Dunne, Timothy; Roberts, Mark J.; and Samuelson, Larry. "The Impact of Plant Failure on Employment Growth in the U.S. Manufacturing Sector." Manuscript. University Park: Pennsylvania State Univ., 1986.

———. "Patterns of Entry and Exit in U.S. Manufacturing Industries." Manuscript. University Park: Pennsylvania State Univ., 1987.

———. "Firm Entry and Post Entry Performance in the U.S. Chemical Industries." Manuscript. University Park: Pennsylvania State Univ., 1988.

———. "Plant Turnover and Gross Employment Flows in the U.S. Manufacturing Sector." *J. Labor Econ.* 7 (January 1989): 48–71.

Ericson, R., and Pakes, Ariel. "An Alternative Theory of Firm and Industry Dynamics." Manuscript. Madison: Univ. Wisconsin, 1989.

Evans, David S. "The Relationship between Firm Growth, Size, and Age: Estimates for 100 Manufacturing Industries." *J. Indus. Econ.* 35 (June 1987): 567–81. (*a*)

———. "Tests of Alternative Theories of Firm Growth." *J.P.E.* 95 (August 1987): 657–74. (*b*)

Hansen, Gary D. "Indivisible Labor and the Business Cycle." *J. Monetary Econ.* 16 (November 1985): 309–27.

Hopenhayn, Hugo A. "Notes on Industry Equilibrium." Manuscript. Stanford, Calif.: Stanford Univ., 1990.

———. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica* 60 (September 1992): 1127–50.

Jovanovic, Boyan. "Selection and the Evolution of Industry." *Econometrica* 50 (May 1982): 649–70.

Kydland, Finn E., and Prescott, Edward C. "Time to Build and Aggregate Fluctuations." *Econometrica* 50 (November 1982): 1345–70.

Lambson, Val. "Industry Evolution with Sunk Costs and Uncertain Market Conditions." Working paper. Madison: Univ. Wisconsin, 1988.

Lazear, Edward P. "Job Security Provisions and Employment." *Q.J.E.* 105 (August 1990): 699–726.

Leonard, Jonathan S. "In the Wrong Place at the Wrong Time: The Extent of Frictional and Structural Unemployment." In *Unemployment and the Structure of Labor Markets*, edited by Kevin Lang and Jonathan S. Leonard. New York: Blackwell, 1987.

Pakes, Ariel, and Ericson, R. "Empirical Implications of Alternative Models of Firm Dynamics." Manuscript. New Haven, Conn.: Yale Univ., 1987.

Rogerson, Richard. "Indivisible Labor, Lotteries, and Equilibrium." *J. Monetary Econ.* 21 (January 1988): 3–16.

Troske, Kenneth. "The Life Cycle of Firms and the Expansion and Contraction of Industries." Manuscript. Chicago: Univ. Chicago, 1989.