Fixed Adjustment Costs, Myopia, and Aggregate Neutrality*

Michael W. L. Elsby
University of Edinburgh

Ryan Michaels
University of Rochester

Preliminary and incomplete

August 29, 2011

Abstract

This paper studies the analytics of a canonical model of fixed adjustment costs in the presence of idiosyncratic productivity shocks. We provide analytical characterizations of the optimal policy function, the aggregate steady state, and the aggregate dynamics implied by the model. These are used to derive a set of approximations to model outcomes in the presence of a small adjustment cost. These reveal that the optimal decisions of firms are approximately myopic, and that a small fixed adjustment cost is neutral with respect to both aggregate steady-state outcomes and aggregate dynamics. A set of numerical illustrations suggests these results are quantitatively relevant for parameterizations commonly used in the literature on employment adjustment.

*We thank Rudi Bachmann for very helpful advice, and participants at the Federal Reserve Bank of New York, and the May 2011 Essex Economics and Music conference. All errors are our own.

E-mail addresses for correspondence: mike.elsby@ed.ac.uk, and ryan.michaels@rochester.edu.

Inaction in microeconomic adjustment is a pervasive phenomenon. A stylized fact of the dynamics of employment, investment and prices is that they exhibit long periods of inaction punctured by bursts of large changes in the number of workers (see, among others, Hamermesh, 1989; Doms and Dunne, 1998; Bils and Klenow, 2005). A leading explanation of this "lumpiness" in microeconomic dynamics is that firms face a fixed cost of adjusting.

In this paper, we study a canonical model of fixed employment adjustment costs. While a substantial literature has analyzed such models numerically, analytical characterizations of optimal labor demand and aggregate employment dynamics have been more difficult to obtain in this environment. This paper seeks to fill this gap.

In section 1, we study the properties of the optimal labor demand policy of an individual firm. In the model, firms face idiosyncratic shocks to labor productivity that induce changes in their desired level of employment. The presence of a fixed adjustment cost leads firms to choose to adjust only infrequently in response to sufficiently large shocks to their productivity. Firms therefore face a potentially complicated employment decision in this environment: They know that the adjustment cost will impede employment adjustment in response to future shocks, and so their current employment decisions will persist into the future with positive probability.

The first result of the paper is to show that these forward-looking considerations are almost irrelevant to the firm in the presence of a small fixed adjustment cost—that is, myopia is nearly optimal. The intuition for this result is one that echoes throughout the paper. We show that a particular symmetry emerges as the adjustment cost becomes small. A firm choosing employment knows that with positive probability it will not adjust employment in the future, and therefore will have to "live with" its current choice. When the adjustment cost is small, however, the positive effect of a marginally higher level of employment in the state of the world where productivity rises is just offset by the negative effect in the state of the world where productivity falls. In expectation, the marginal effect of current employment decisions on future profits is approximately zero.

We examine the quantitative relevance of this approximation by analyzing numerically a calibration of the model that corresponds to estimates obtained in the recent empirical literature on employment adjustment (Cooper, Haltiwanger and Willis, 2005, 2007; Bloom, 2009). Although some small differences can be discerned between the optimal forward-looking labor demand policy and its myopic counterpart, the approximation is in general quite accurate.

This result is valuable for two reasons. It provides a simple insight into the form of optimal labor demand in what otherwise might be seen as a complicated economic environment. In addition, we will see that the simplicity of the myopic policy will faciliate the analysis of aggregate equilibrium, to which we now turn.

In section 2, we take on the task of aggregating this behavior to the macroeconomic level. The aggregate implications of lumpy microeconomic adjustment also are not obvious. Since firms follow a highly nonlinear optimal labor demand policy, and face heterogeneous idiosyncratic productivities, a representative firm interpretation of the model is not available. The second contribution of the paper is to show how it is possible to characterize the steady-state distribution of employment across firms by applying a simple mass-balance approach. We show that the firm-size distribution satisfies a recursion with a simple interpretation. It reveals that the steady-state density at some level of employment is equal to its frictionless counterpart multiplied by the ratio of the probability of adjusting to that level of employment to the probability of adjusting away from that level.

This simple interpretation in turn motivates an additional approximation result. We show that, again for a sufficiently small adjustment cost, the steady-state distribution of employment coincides with its frictionless counterpart—that is, that the adjustment cost is approximately neutral with respect to aggregate outcomes. This in turn implies that the aggregate demand for labor also is approximately invariant to the adjustment cost, since it is implied by the mean of the firm-size distribution.

The key to understanding this result again can be traced to a symmetry of the model implied by the myopic approximation in section 1. The presence of an adjustment cost reduces the probability of adjusting to an employment level, but it also reduces the probability of adjusting away from that level. In the neighborhood of a small adjustment cost, these two opposing effects cancel exactly, and the adjustment cost is neutral with respect to steady-state aggregate outcomes.

To examine the quantitative relevance of the approximate steady-state invariance result, we return to the numerical model analyzed in section 1. Consistent with the approximation result, the aggregate steady-state labor demand schedule implied by the model with a fixed adjustment cost lies close to that implied by the frictionless model. Similarly, a set of simulated aggregate moments—the mean and standard deviation of employment across firms, and the probability of adjusting—are very similar to their frictionless counterparts.

In our final set of results, in section 3 we study the aggregate dynamics of the model in the presence of aggregate shocks. We show how it is possible to use the mass balance approach that underlies the steady-state results to characterize analytically the dynamic evolution of the distribution of employment across firms, and thereby of aggregate employment also. This in turn allows us to show that the same symmetry that underlies approximate aggregate steady-state neutrality also holds along the dynamic transition path. Mirroring the intuition for the steady-state results, a small fixed cost reduces the outflow of mass from a given employment level, but also reduces the mass of firms which find it optimal to adjust to that employment level. For small frictions, these two forces offset and leave the distribution approximately equal to its frictionless counterpart period by period. Thus, aggregate employment is predicted to display near-jump dynamics, even in the presence of a fixed adjustment cost.

We find that this dynamic invariance result is borne out in numerical simulations of the dynamic response of aggregate outcomes based on the calibration introduced in section 1. Impulse responses of aggregate employment to an innovation to aggregate productivity reveal responses that lie very close to the frictionless analogue.

In the concluding sections of the paper, we discuss how our results dovetail with the large literature on adjustment costs. We highlight an interesting feature of our approximate aggregate invariance results, namely that it does not rely on general equilibrium adjustment of wages. This contrasts with Kahn and Thomas' (2008) recent influential work on investment adjustment costs, who emphasize these general equilibrium forces. In addition, we also revisit two key papers in the recent literature on dynamic labor demand. In simulations of a model without productive heterogeneity, King and Thomas (2006) find that the response of aggregate employment to aggregate productivity displays a slight hump-shape, in contrast to the near-jump dynamics we find. Similarly, Bachmann (2009) observes a similar result in a model with productive heterogeneity, but with a much lower average adjustment rate. We suggest that this highlights the circumstances in which neutrality can be expected to arise, namely when the adjustment cost is sufficiently small relative to the volatility of idiosyncratic shocks. While our quantitative results suggest recent estimates fall within this range, we argue that future empirical work needs to focus on obtaining robust estimates of these parameters in order to infer the role of lumpy microeconomic adjustment in aggregate employment dynamics.

1 The Firm's Problem

We consider a canonical model of fixed employment adjustment costs. Time is discrete. Firms use labor, n, to produce output according to the production function, $y = pxF(n)$, where p represents the state of aggregate labor demand; x represents shocks that are idiosyncratic to an individual firm; and the function F is increasing and concave, $F_n > 0$ and $F_{nn} < 0$. We assume that the evolution of the idiosyncratic shocks is described by a continuously differentiable distribution function, $G(x'|x)$.

We begin by analyzing the model in steady state, so for now we treat p as fixed. At the beginning of a period, a firm observes the realization of its idiosyncratic shock, x. At the same time, an exogenous fraction δ of its workforce separates. Given this, they then make their employment decision. If the firm chooses to adjust the size of its workforce, it incurs a fixed adjustment cost, denoted C.

It is common in the literature to scale the adjustment cost so that very productive firms do not "outgrow" it. When coupled with other assumptions on the stochastic process of x, this scaling can imply a form of homogeneity that simplifies the dynamic problem of a firm (see, for example, Caballero and Engel, 1999, and Gertler and Leahy, 2008). To emphasize that the following results do not rely on this homogeneity, we focus initially on a pure lump-sum adjustment cost. Towards the end of the section, however, we show that our results nonetheless extend to this case.

It follows that we can characterize the expected present discounted value of a firm's profits recursively as:[1]

$$\Pi(\tilde{n}_{-1}, x) \equiv \max_n \left\{ pxF(n) - wn - C\mathbf{1}^{\Delta} + \beta \int \Pi(\tilde{n}, x') \, dG(x'|x) \right\}, \tag{1}$$

where $\tilde{n}_{-1} \equiv (1-\delta)\, n_{-1}$ denotes employment carried into the period, and $\mathbf{1}^{\Delta} \equiv \mathbf{1}[n \neq \tilde{n}_{-1}]$ is an indicator that equals one if the firm adjusts and zero otherwise. The wage w is determined in a competitive labor market, and is taken as exogenous from the firm's perspective.[2]

For the analysis that follows, it is helpful to recast the firm's problem in equation (1) into two related underlying Bellman equations. In particular, the value of adjusting (gross of the adjustment cost), $\Pi^{\Delta}(x)$, and the value of not adjusting, $\Pi^{0}(\tilde{n}_{-1}, x)$, are given by

$$\begin{aligned} \Pi^{\Delta}(x) &\equiv \max_n \left\{ pxF(n) - wn + \beta \int \Pi(\tilde{n}, x') \, dG(x'|x) \right\}, \text{ and} && (2) \\ \Pi^{0}(\tilde{n}_{-1}, x) &\equiv pxF(\tilde{n}_{-1}) - w\tilde{n}_{-1} + \beta \int \Pi((1-\delta)\,\tilde{n}_{-1}, x') \, dG(x'|x). && (3) \end{aligned}$$

[1] We adopt the convention of denoting lagged values with a subscript, $_{-1}$, and forward values with a prime, $'$.

[2] The law of one wage can be supported by assuming that workers are perfectly mobile (and thus may seek new job opportunities at any point in time). One might wonder whether a worker is able to hold up a firm and demand a higher wage, since her departure would appear to force the firm to pay an adjustment cost. This strategy is feasible only if the present discounted value of work among firms that are hiring exceeds that at her present firm. We will see in equation (**), however, that firms hire until the value of the marginal worker is zero. Thus, marginal hires yield no rents. It follows that all workers at all firms are paid just enough to make them indifferent between work and non-work. Moreover, if firms cannot commit, this contract must be implemented as a period-by-period, economy-wide wage.

Clearly, the value of the firm $\Pi(\tilde{n}_{-1}, x)$ is simply the upper envelope of these two regimes,

$$\Pi(\tilde{n}_{-1}, x) = \max\left\{\Pi^{\Delta}(x) - C, \Pi^{0}(\tilde{n}_{-1}, x)\right\}. \tag{4}$$

In keeping with the literature on fixed adjustment costs, we assume that the optimal labor demand policy takes an Ss form.[3] Figure 1 illustrates such a policy. It is characterized by three functions, $L(n) < X(n) < U(n)$. In the event that the firm chooses to adjust away from $\tilde{n}_{-1}$, optimal employment is determined by a "reset" function $X(n)$ which satisfies the first-order condition

$$pX(n)F_n(n) - w + \beta(1-\delta)\int \Pi_1(\tilde{n}, x')\, dG(x'|X(n)) \equiv 0. \tag{5}$$

Due to the adjustment cost, however, the firm will not always choose to adjust: It will decide to adjust only if the value of adjusting, net of the adjustment cost, $\Pi^{\Delta}(x) - C$, exceeds the value of not adjusting, $\Pi^{0}(\tilde{n}_{-1}, x)$. This aspect of the firm's decision rule is characterized by two adjustment "triggers," $L(\tilde{n}_{-1})$ and $U(\tilde{n}_{-1})$. For sufficiently bad realizations of the idiosyncratic shock, $x < L(\tilde{n}_{-1})$, the firm will shed workers. For sufficiently good shocks, $x > U(\tilde{n}_{-1})$, the firm will hire workers. For intermediate values of $x \in [L(\tilde{n}_{-1}), U(\tilde{n}_{-1})]$, the firm will neither hire nor fire, and $n = \tilde{n}_{-1}$. Thus, the adjustment triggers trace out the locus of points for which the firm is indifferent between adjusting and not adjusting. It follows that the triggers satisfy the value-matching conditions

$$\Pi^{\Delta}(L(\tilde{n}_{-1})) - C = \Pi^{0}(\tilde{n}_{-1}, L(\tilde{n}_{-1})), \text{ and } \Pi^{\Delta}(U(\tilde{n}_{-1})) - C = \Pi^{0}(\tilde{n}_{-1}, U(\tilde{n}_{-1})). \tag{6}$$

1.1 The Near-Optimality of Myopia

In this subsection, we show how it is possible to gain further insights into the form of the optimal labor demand policy by taking analytical approximations to the firm's problem. We show that, from a theoretical perspective, a firm's incentive to consider the future consequences of their employment decisions is second order in the adjustment cost—that is, myopia is nearly-optimal:

Proposition 1 *In the presence of a small fixed adjustment cost, the optimal labor demand*

[3] It is well-known that it is difficult to prove the optimality of the Ss policy in settings outside the canonical Brownian model (Harrison, Sellke and Taylor, 1983). In what follows we show that, in the neighborhood of a small adjustment cost, the optimal labor demand policy is well-approximated by its myopic counterpart. A useful implication of this result is that, since we know the optimal myopic policy takes the Ss form, we also know that the Ss policy is approximately optimal for the context of the results in this paper.

policy of the firm coincides with the myopic ($\beta = 0$) solution.

To understand this result, note that the implications of current employment decisions for future profits are summarized by the forward value in equation (1), $\int \Pi(\tilde{n}, x') \, dG(x'|x)$. Proposition 1 states that the latter is independent of current employment, n, for small adjustment costs, $\int \Pi(\tilde{n}, x') \, dG(x'|x) \approx 0$ for $C \approx 0$.

To see why this is so, observe first that one may partition the forward value into adjustment and inaction regimes:

$$\Pi(\tilde{n}, x') = \begin{cases} \Pi^{\Delta}(x') - C & \text{if} \quad x' \notin [L(\tilde{n}), U(\tilde{n})], \\ \Pi^{0}(\tilde{n}, x') & \text{if} \quad x' \in [L(\tilde{n}), U(\tilde{n})]. \end{cases} \tag{7}$$

The relationship between future profits and current employment therefore will resemble that illustrated in the two panels of Figure 2. Panel A fixes next period's idiosyncratic productivity, and traces out future profits as a function of current employment. The best possible outcome for the firm is realized when employment carried into next period, $\tilde{n}$, turns out to be equal to the firm's desired employment level, $X^{-1}(x')$. Away from that point, future profits fall in both directions. Importantly, however, if current employment is either so low, or so high, that the firm will adjust next period, the present choice of n has no effect on the future value of the firm. This implies that the marginal effect of current employment decisions on future profits is given by $\int_{L(\tilde{n})}^{U(\tilde{n})} \Pi_1^0(\tilde{n}, x') \, dG(x'|x)$, since the trigger strategy imposes that it is zero outside of the inaction region.

As future productivity x' rises, firms' desired employment also rises, and so the function in Panel A shifts upward and to the right. It follows that the marginal effect of current employment on future profits, $\Pi_1(\tilde{n}, x')$, will be as depicted in Panel B of Figure 2. For values of future productivity that induce adjustment ($x' \notin [L(\tilde{n}), U(\tilde{n})]$), the marginal effect is zero. Within the inaction band ($x' \in [L(\tilde{n}), U(\tilde{n})]$), however, $\Pi_1(\tilde{n}, x')$ is increasing, crossing zero at the point where employment brought into the period is optimal, $x' = X(\tilde{n})$.

The key to understanding why myopia is nearly-optimal in this enviroment is to note that a certain symmetry emerges as C becomes small. Specifically, in the neighborhood of a small adjustment cost, one can approximate the gross return to adjusting, $\Delta(\tilde{n}, x') \equiv \Pi^{\Delta}(x') - C - \Pi^{0}(\tilde{n}.x')$ as[4]

$$\Delta(\tilde{n}, x') \approx \frac{1}{2} \Delta_{xx}(\tilde{n}, X(\tilde{n})) (x' - X(\tilde{n}))^2. \tag{8}$$

[4]By virtue of the optimality of $\tilde{n}$ when $x = X(\tilde{n})$, the first and second terms in the expansion are identically equal to zero: $\Delta(\tilde{n}, X(\tilde{n})) \equiv 0$ and $\Delta_x(\tilde{n}, X(\tilde{n})) \equiv 0$.

It follows from the value-matching condition (6) that the adjustment triggers $L(\tilde{n})$ and $U(\tilde{n})$ become symmetric around the reset point $X(\tilde{n})$:

$$L(\tilde{n}) \approx X(\tilde{n}) - \gamma(\tilde{n})\sqrt{C}, \text{ and } U(\tilde{n}) \approx X(\tilde{n}) + \gamma(\tilde{n})\sqrt{C}, \tag{9}$$

where $\gamma(\tilde{n}) \equiv \sqrt{2/\Delta_{xx}(\tilde{n}, X(\tilde{n}))}$.

A related argument implies that the forward value $\Pi(\tilde{n}, x')$ in Figure 2A also is approximately quadratic in $\tilde{n}$, and hence also symmetric, in the neighborhood of a small adjustment cost. Thus, the marginal effect of current employment on future profits, $\Pi_1(\tilde{n}, x')$, becomes approximately linear within the inaction region, as depicted in Figure 2B. This observation, together with equation (9), implies that the marginal future value of current employment becomes approximately $\int_{X(\tilde{n})-\gamma(\tilde{n})\sqrt{C}}^{X(\tilde{n})+\gamma(\tilde{n})\sqrt{C}} [x' - X(\tilde{n})]\, dG(x'|x)$.

Figure 2B therefore reveals why this symmetry property of the model lies at the heart of Proposition 1. As the adjustment cost becomes small, the positive effect of carrying a marginally higher level of employment into next period when future productivity is relatively high ($x' > X(\tilde{n})$) is offset exactly by the negative effect in the state of the world where future productivity is relatively low ($x' < X(\tilde{n})$). It follows, then, that the *expected* marginal effect of current employment on future profits, $\int \Pi_1(\tilde{n}, x')\, dG(x'|x)$, is approximately zero—myopia is nearly optimal.

Note that this implies the forward value in the firm's problem (1) is approximately independent of the adjustment cost for *any* level of employment carried into next period $\tilde{n}$. Thus, both the reset function defined in (5), as well as the adjustment triggers in (6), will coincide with their myopic counterparts in the neighborhood of a small adjustment cost.

1.2 Numerical Illustration

In Proposition 1, we established that myopia is nearly optimal, so long as the adjustment cost is sufficiently small. In this subsection, we assess the quantitative relevance of the result. We do this by analyzing numerically a parameterization of the model that corresponds to estimates obtained in recent literature.

In order to mirror the framework used in this literature, we need to place a little more structure on the problem. First, we adopt the widespread assumption that the production function is of the Cobb-Douglas form, $F(n) = n^{\alpha}$, with $\alpha < 1$. Second, as noted above, it is common to scale the adjustment cost by a measure of firm size so that larger, more productive firms do not "outgrow" the friction. In keeping with this, we adopt a specification similar to that assumed in Gertler and Leahy (2008) in which the adjustment cost is proportional to the firm's (frictionless)[5] revenue, which we denote $R(x)$. Thus, $C(x) = cR(x)$ for some $c > 0$. The following corollary confirms that the myopic approximation extends to the version of the standard model with size-dependent frictions:[6]

Corollary 1 *Suppose a firm uses a production function $F(n) = n^\alpha$ and faces a fixed adjustment cost of the form, $C(x) = cR(x)$. If c is sufficiently small, the myopic approximation of Proposition 1 continues to hold.*

Finally, we specialize the stochastic process governing idiosyncratic shocks, $G(x'|x)$, to the common assumption of a geometric AR(1),

$$\log x' = \mu_x + \rho_x \log x + \varepsilon'_x, \text{ where } \varepsilon'_x \sim N\left(0, \sigma_x^2\right). \tag{10}$$

Given these assumptions, we solve the model numerically in order to determine the quantitative accuracy of the myopic approximation.

The baseline calibration we analyze is summarized in Table 1. The numerical model is cast at a quarterly frequency. Our calibration of the magnitude of the adjustment cost is based on estimates reported in Cooper, Haltiwanger and Willis (2005, 2007) and Bloom (2009). Cooper, Haltiwanger and Willis (2005) estimate a model similar to the one described above using plant-level data from the Census' Longitudinal Research Database. They estimate a fixed cost of adjustment equal to approximately 8 percent of quarterly revenue (see row "Disrupt"in their Table 3A).[7] Using Compustat data, Bloom (2009) finds nearly the same result (see column "All"in his Table 3). Based on this, we calibrate the adjustment cost parameter c to equal 0.08.

There is less consensus over the parameters of the process of idiosyncratic shocks in equation (10). Cooper, Haltiwanger and Willis (2005) obtain estimates of $\rho_x = 0.39$ and

[5] That the cost is assumed proportional to frictionless revenue, as opposed to revenue in the presence of the adjustment cost, is not consequential. For example, an alternative interpretation is that the adjustment cost is approximately proportional to the revenue of an adjusting firm. To see why, note that this will be exactly the case for an adjusting firm which implements the myopic labor demand policy (namely, $n = [\alpha p x / w]^{1/(1-\alpha)}$), which we will see is a good approximation to the forward-looking policy.

[6] The only difference that arises is that the symmetry property underlying Proposition 1 holds only in the limit in the case of a size-dependent friction. This nonetheless turns out to be sufficient to establish first-order invariance of the forward value to the adjustment cost.

[7] The model Cooper et al. estimate is a little more elaborate than the one considered in this section. First, Cooper et al. estimate a "disruption" cost that is proportional, not to current revenue, but to the plant's revenue if it adjusts. Second, Cooper et al. allow for intensive margin adjustments of hours worked. We have solved numerically a version of the model with both of these features, and found that the myopic approximation continues to hold very well. A final difference is that Cooper et al.'s model allows for convex adjustment costs. We abstract from the latter to focus more precisely on the effect of non-convex frictions.

$\sigma_x = 0.5$. In a later paper, Cooper, Haltiwanger and Willis (2007) estimate these parameters within the context of a search-and-matching model using, in part, monthly establishment-level data from the Job Openings and Labor Turnover Survey. Most of their estimates of ρ_x are, when converted to a quarterly frequency, much smaller than 0.39. Moreover, their estimates of σ_x (again, converted to a quarterly frequency) are notably lower, near 0.2.[8] A value of $\rho_x < 0.39$ runs aggressively against evidence from other data sources (see Foster, Haltiwanger, and Syverson, 2008). For the purposes of the baseline calibration, we set $\rho_x = 0.4$ at a quarterly frequency, and split the difference between Cooper, Haltiwanger and Willis (2005) and (2007) and set $\sigma_x = 0.35$. This choice is comparable to the calibration in Bachmann's (2009) analysis of non-convex adjustment costs.[9]

The remaining four parameters are more straightforward to calibrate. The returns to scale parameter α is set equal to 0.64, based on the estimates in Cooper, Haltiwanger and Willis (2005). This also is similar to the value assumed by King and Thomas (2006) in their analysis of a related fixed-cost model. Second, the discount factor β is set to 0.99, which is the conventional choice for a quarterly model. Third, we set the rate of exogenous worker attrition δ equal to 0.06, which corresponds to the average quarterly quit rate in the Job Openings and Labor Turnover Survey. Finally, since general equilibrium considerations are not relevant for the present discussion, we the (constant) wage rate is fixed such that average employer size is 20. This is roughly consistent with data from County Business Patterns.

Given this calibration, we solve the dynamic model numerically. Based on Proposition 1 and its corollary, we solve the model via policy function iteration, using the myopic policy rule as the initial condition. The details of this procedure are discussed in Judd (1998).

Figure 3 graphs the optimal labor demand policies in the myopic and forward-looking versions of the model. The three functions, $L(n)$, $X(n)$, and $U(n)$, are illustrated in bold for the forward-looking policy, and as shaded lines for the myopic case. To provide two alternative perspectives of any deviations between the policies, they are plotted on both levels (Panel A) and logarithmic (Panel B) scales. Taken as a whole, Figure 3 suggests that the myopic policy is quite similar to the optimal forward-looking policy, as suggested by

[8] The typical estimate of ρ_x in Cooper, Haltiwanger and Willis (2007) is roughly 0.4 at a *monthly* frequency. Therefore, the implied degree of quarterly persistence, $0.4^3 = 0.064$, is very low. Likewise, their estimates of the variance of the monthly innovation is in the neighborhood $\sigma_x^2 \approx 0.2^2 = 0.04$. Its counterpart in a quarterly model is therefore $\sqrt{(0.4^4 + 0.4^2 + 1)\, 0.04} \approx 0.22$.

[9] We have begun to assess the sensitivity of our results to alternative calibrations of $\{\rho_x, \sigma_x\}$, and these will be folded into the next draft of this paper. In short, variations in ρ_x, given σ_x, have relatively minimal effects on the impulse responses reported in section 3. Reductions in σ_x, on the other hand, can cause more noticeable departures from the frictionless outcome (in particular, the impulse response displays a slight hump shape), but the differences are not dramatic.

Proposition 1.

There are differences, though. Specifically, at lower values of idiosyncratic productivity x, we observe that a forward-looking firm will choose a slightly higher level employment than a myopic firm, and vice versa at relatively high values of x (although the latter is barely perceptible). Intuitively, a forward-looking, low (high) productivity firm adjusting this period will set employment somewhat higher (lower) than its myopic counterpart in anticipation of reversion toward the mean.

Figure 3 also provides a sense of the magnitude of these deviations. For example, the absolute difference in optimal employment between the two policies in Figure 3A averages a little over one worker. However, as Figure 3B reveals, such a deviation is larger in percentage terms for smaller firms.

Another way to quantify the similarity of the myopic and forward-looking policy rules is to simulate the employment paths implied by each when a firm receives the same productivity sequence. Figure 4 presents a scatter plot of the two series. The employment paths lie very close to one another: The slope is 1.03 and the R^2 exceeds 0.97.

1.3 Discussion

Our finding that myopic behavior is nearly optimal in the presence of a small adjustment friction is significant both in the context of previous literature, and for the results that follows in remainder of this paper.

At a basic level the result is surprising because a quintessential feature of the large literature on adjustment costs is notion that forward-looking behavior is important. Proposition 1 suggests that these considerations are not a first-order issue in the presence of a small fixed adjustment cost. There are precedents for this result, however. In their Simplification Theorem, Gertler and Leahy (2008) show that, in the presence of a second-order small adjustment cost, and idiosyncratic shocks that evolve according to a geometric random walk with uniform innovations, current decisions affect future profits only through a term that is third-order small. Thus, the near-optimality of myopia is implicit in their analysis. What Proposition 1 suggests is that this result extends beyond the special case in which shocks follow a log-uniform random walk. The only restriction on the evolution of idiosyncratic shocks required in Proposition 1 is that the density $g(x'|x)$ be continuous in x'.[10]

[10] We should note that Gertler and Leahy consider a problem that is slightly more elaborate than the one we examine. They assume that idiosyncratic productivity follows a compound Poisson process whereby, if a firm receives a shock, its new productivity is drawn according to a geometric random walk with uniform innovations. *Conditional on receipt* of the shock, Gertler and Leahy prove the near-optimality of myopia.

The myopic approximation is also useful from a practical perspective. It has the virtue of being very easy to solve for. In particular, Proposition 2 reveals that all one need do is find the roots of a simple nonlinear equation.

Proposition 2 *If a firm is myopic ($\beta = 0$), its optimal labor demand policy takes an* Ss *form. Moreover, if the production technology is $F(n) = n^{\alpha}$, and if the adjustment cost is $C(x) = cR(x)$, then the policy is loglinear:*

$$L(n_{-1}) = \mathcal{L} n_{-1}^{1-\alpha},\ X(n_{-1}) = \mathcal{X} n_{-1}^{1-\alpha},\ \textit{and}\ U(n_{-1}) = \mathcal{U} n_{-1}^{1-\alpha}, \tag{11}$$

where $\mathcal{L}$ and $\mathcal{U}$ are the roots $\mathcal{Z}$ of $(1-\alpha-c)(\mathcal{Z}/\mathcal{X})^{\frac{1}{1-\alpha}} = (\mathcal{Z}/\mathcal{X}) - \alpha$, and $\mathcal{X}$ coincides with its frictionless counterpart $\mathcal{X} \equiv w/\alpha p$.

The majority of the literature on fixed adjustment costs has applied numerical methods, such as value function iteration, to solve for optimal policies. This can be a time-consuming process. By using the myopic policy as an initial conjecture, computation time can be reduced considerably.

It can be also be difficult to generalize from specific parameterizations using numerical methods. Proposition 1 suggests that the myopic solution can be used to infer the (approximate) effect of exogenous parameters on a firm's policy function. In what follows, we use the simplicity of the myopic policy to infer certain properties of the aggregate implications of fixed adjustment costs, in the knowledge that the myopic rule in turn acts as a good approximation to the optimal labor demand policy. This is what we do in the next section.

2 Aggregation and Steady State Equilibrium

In this section, we address the aggregate implications of lumpy adjustment at the microeconomic level. Aggregation in this context non-trivial. Each individual firm's labor demand policy depends in a highly nonlinear fashion on both their individual lagged employment n_{-1}, as well the realization of their idiosyncratic shock x. Since firms are heterogeneous in these state variables, it follows that there is no simple representative firm interpretion of the model.

In what follows, we develop an analytical approach that allows one to aggregate microeconomic behavior to the macroeconomic level when firms face a fixed cost of adjusting the

The probability that the firm does not draw a new shock acts like a Calvo probability of not adjusting. It therefore serves as an exogenous restriction on future adjustment. For this reason, the optimal policy in their model is forward-looking, but not as a result of the fixed adjustment cost.

size of their workforces. In doing so, we obtain a characterization of the steady-state distribution of employment across firms, the mean of which is aggregate employment. This allows us to infer aggregate labor demand, and thereby steady-state equilibrium.

By deriving analytical expressions for these aggregates, we are in turn able to derive approximations to these outcomes in the neighborhood of a small adjustment cost, mirroring the analysis of section 1.1. These approximations turn out to be particularly simple for the myopic case. In particular we show that, in the presence of sufficiently small adjustment frictions, the steady-state distribution of employment under myopia is approximately identical to that in the frictionless model. Since we know that the myopic policy serves as a good approximation to its forward-looking counterpart, a corollary is that this approximate neutrality also extends to the more general forward-looking case—i.e. that is, aggregate employment is approximately invariant to a small fixed adjustment cost.

We then examine the quantitative relevance of this result by analyzing numerically the aggregate implications of the baseline calibration in section 2.3. The results of this exercise reveal that the approximation holds up quite well.

2.1 Analytics of Aggregate Labor Demand

In order to infer the aggregate demand for labor, and hence solve for steady-state aggregate equilibrium, we first solve for a related object, the steady-state distribution of employment across firms, denoted $H(n)$. We construct this distribution by applying a mass-balance approach—that is, by setting the inflow into the mass $H(n)$ equal to the outflows from that mass. While the use of mass balance is a relatively standard technique in many domains of economics, the application to the fixed adjustment cost model is non-trivial and, to the best of our knowledge, new to the literature. Proposition 3 summarizes the results of this exercise.[11][12]

Proposition 3 *The steady-state density of employment across firms is given by*

$$h\left(n\right) = h^{*}\left(n\right) \frac{1 - \mathcal{H}\left[L^{-1}X\left(n\right), X\left(n\right)\right] + \mathcal{H}\left[U^{-1}X\left(n\right), X\left(n\right)\right]}{1 - \mathcal{G}\left[U\left(n\right), n\right] + \mathcal{G}\left[L\left(n\right), n\right]}, \tag{12}$$

where $h^{*}(n)$ is the frictionless density of employment, $\mathcal{H}\left(\nu, \xi\right) \equiv \Pr\left[n_{-1} \leq \nu | x = \xi\right]$, and $\mathcal{G}\left(\xi, \nu\right) \equiv \Pr\left[x \leq \xi | n_{-1} = \nu\right]$.

[11] For notational simplicity, in what follows we denote composite functions, such as $L^{-1}\left[X\left(n\right)\right]$ as $L^{-1}X\left(n\right)$, and so on.

[12] The statement of Proposition ** here holds for the case in which there is no exogenous attrition, $\delta = 0$. The extended proof that allows for $\delta > 0$ is in progress.

To see the significance of this result, it is instructive to consider the special case in which the idiosyncratic shocks faced by the firm are i.i.d. across time. In that case, a firm's lagged employment n_{-1} provides no information on the distribution of its current productivity x, and vice versa. It follows that the steady-state density of employment in equation (12) takes the simpler form

$$h(n) = h^*(n) \frac{1 - H\left[L^{-1}X(n)\right] + H\left[U^{-1}X(n)\right]}{1 - G\left[U(n)\right] + G\left[L(n)\right]}, \tag{13}$$

Thus, the density of employment $h(n)$ is defined recursively—it depends upon its integral, the distribution function of employment $H(\cdot)$ evaluated at two different points in its domain in the form of a difference-differential equation.

The special case in which idiosyncratic shocks are i.i.d. also clarifies the intuition behind Proposition 3. The result states that the steady-state density of employment $h(n)$ is proportional to its frictionless counterpart $h^*(n)$. To understand the factor of proportionality, consider first the denominator in equation (13), $1 - G\left[U(n)\right] + G\left[L(n)\right]$. Inspection of the illustration of the policy function depicted in Figure 1 reveals that this is simply the probability that a firm with employment level n adjusts away from n.

The numerator of equation (13) is more difficult to interpret. Consider a firm that draws an idiosyncratic productivity level $x = X(n)$. Absent an adjustment cost, such a firm would adjust to an employment level of n. In the presence of an adjustment cost, however, Figure 1 reveals that the firm's decision to adjust or not will depend on the level of employment the firm has inherited from the past. Firm whose initial employment is either relatively low ($n_{-1} < U^{-1}X(n)$) or relatively high ($n_{-1} > L^{-1}X(n)$) will adjust to n. Thus, the numerator of equation (13) is simply the probability that a firm with productivity $x = X(n)$ adjusts to n.

Returning to the general result reported in Proposition 3, it is clear that the same interpretation extends to the case with persistent idiosyncratic shocks in equation (12). To summarize, then, Proposition 3 may be restated simply as

$$h(n) = h^*(n) \frac{\Pr(\text{adjust to } n)}{\Pr(\text{adjust from } n)}. \tag{14}$$

This is a very intuitive result. Put simply, if a firm is more likely to adjust to n than it is to adjust away from n, then the distribution of employment will accumulate more mass at n.

Proposition 3 also has a number of other useful features. First, it provides a clear link from microeconomic behavior in the model to the aggregate outcomes implied by that behavior. Specifically, once one knows the optimal labor demand policy used by individual firms, as summarized by the functions $L(n) < X(n) < U(n)$, we can infer the implied distribution of employment in the aggregate.

This in turn allows us to trace out the aggregate demand for labor, which is implied by the mean of $h(n)$. In particular, the optimal labor demand policy analyzed in section 1.1 is defined implicitly for a given level of the market wage w. Thus, the solution for the steady-state distribution of employment in (12) also is indexed implicitly by w. Making that dependence explicit, one can express the aggregate demand for labor as

$$N^d(w) = \int nh(n; w)\, dn. \tag{15}$$

Thus, when combined with a model of labor supply, our aggregation result allows one to determine aggregate steady-state equilibrium employment and wages.

2.2 The Near-Invariance of Steady-State Aggregate Outcomes

In this subsection, we return to the case in which the adjustment friction is small, $C \approx 0$, and consider its implied aggregate effects. We show that the distribution of employment across firms is second order in the adjustment cost. That is, the adjustment cost is approximately neutral with respect to aggregate steady-state employment outcomes. We derive this result in two stages. First, we demonstrate that the distribution of employment is invariant to a small fixed adjustment cost in the case in which firms are myopic (that is, $\beta = 0$). The subsequent corollary builds on Proposition 1 to argue that this invariance result extends to the forward-looking case ($\beta > 0$).

Proposition 4 *If firms are myopic ($\beta = 0$) and face a small fixed adjustment cost ($C \approx 0$), the steady-state distribution of employment across firms $h(n)$, and therefore aggregate labor demand $N^d(w)$, coincide with their frictionless counterparts.*

The intuition for this result can be gleaned by returning to equation (14) above. One feature of (14) is particularly instructive in the present context. Since the existence of an adjustment cost reduces both the probability of adjusting from a given employment level, as well as the probability of adjusting to that level, the impact on the distribution of employment, and hence on aggregate employment, is not obvious *a priori*.

What Proposition 4 suggests is that, in the neighborhood of a small adjustment cost, these two opposing effects cancel exactly, and the adjustment cost is neutral with respect to steady-state aggregate outcomes. That is, the presence of a small fixed adjustment cost reduces the probability of adjusting to an employment level by approximately as much as it reduces the probability of adjusting away from that level.

Mirroring the near-myopia result in section 1, this neutrality result stems from a symmetry that emerges as C becomes small. To see intuitively how this symmetry plays out, the special case in which idiosyncratic shocks are i.i.d. over time is again useful. In that case, the steady-state distribution of employment is given by equation (13). This reveals that the distribution of employment will coincide with its frictionless counterpart, $H(n) = H^*(n) = G[X(n)]$, if the adjustment triggers satisfy the following symmetry property: $U^{-1}X(n) = X^{-1}L(n)$, and $L^{-1}X(n) = X^{-1}U(n)$.

Figure 5 illustrates why this is exactly the symmetry that arises in the presence of a small adjustment cost. From equation (9), we know that the adjustment triggers $U(n)$ and $L(n)$ are approximately symmetric around the reset function $X(n)$. Inverting these, it follows that

$$\begin{aligned} U^{-1}X(n) &\approx X^{-1}\left(X(n) - \gamma(n)\sqrt{C}\right) \approx X^{-1}L(n), \text{ and} \\ L^{-1}X(n) &\approx X^{-1}\left(X(n) + \gamma(n)\sqrt{C}\right) \approx X^{-1}U(n), \end{aligned} \tag{16}$$

as shown in Figure 5. The proof of Proposition 4 formalizes this intuition and shows how it extends to the case of persistent idiosyncratic shocks.

We now wish to use Proposition 4 to make a statement about the steady-state distribution of employment under a forward-looking policy. To do so, recall from Proposition 1 that the optimal labor demand policy coincides with the myopic policy for small C. Since the distribution of employment across firms is determined solely by the policy function (Proposition 3), if the myopic policy induces a distribution of employment that is approximately the same as that in the frictionless model, so does the forward-looking policy rule.

Corollary 2 *If firms are forward-looking ($\beta > 0$), Proposition 4 continues to hold.*

2.3 Numerical Illustration

We now return to the calibration of section 1.2 to assess the quantitative accuracy of the approximation result in Proposition 4. Before we proceed, we first confirm that Proposition 4 extends to the model with a size-dependent friction that we use in our numerical simulations. The following Corollary confirms this.

Corollary 3 *Suppose a firm uses a production function $F(n) = n^\alpha$ and faces a fixed adjustment cost of the form, $C(x) = cR(x)$. If c is sufficiently small, the approximate aggregate* *neutrality result in Proposition 4 continues to hold.*

The intuition here is that, in the place of the arithmetic symmetry of the adjustment triggers in equation (9), in the presence of a size-dependent friction the triggers exhibit an approximate geometric symmetry that is inherited from the loglinearity of the adjustment cost. The Corollary demonstrates that this approximate geometric symmetry shares the property that $U^{-1}X(n) \approx X^{-1}L(n)$ and $L^{-1}X(n) \approx X^{-1}U(n)$ in Figure 5, and thus approximate aggregate steady-state neutrality is preserved.

Our quantitative investigation of the relevance of this approximation compares a set of steady-state outcomes implied by the baseline calibration summarized in Table 1. We examine three versions of the model: the forward-looking model with frictions ($\beta > 0$, $C > 0$), its myopic ($\beta = 0$) counterpart, and the frictionless model ($C = 0$). To begin, Table 2 reports average firm size and a measure of dispersion of the firm size distribution (specifically, the standard deviation of $\log n$). Mean employment in the frictionless and forward-looking models are within 3.5 percent of one another, but the dispersion is more noticeably attenuated under a forward-looking policy. Comparing the myopic and frictionless outcomes, the results are reversed, in a sense: aggregate employment under myopia is almost 9 percent below the frictionless outcome, but the degree of dispersion in the two models is virtually identical.

We suspect that the difference with regard to the mean reflects the absence of precise symmetry of the triggers about the return point, $X(n)$, as discussed in the proofs of Propositions 1 and 4. To see why this happens, recall that the friction takes the form $C(x) = cR(x)$. Importantly, $R(x)$ is convex in x (specifically, $R(x)$ is proportional to $x^{\frac{1}{1-\alpha}}$). As a result, if c is sufficiently large, the adjustment cost depletes more of the gain from adjusting to a *higher* x than it does when adjusting to a *lower* x. This is perhaps why, in Figure 3, productivity must rise by a relatively large amount to induce a firm to hire, and thus why the upper trigger is further from the return point than is the lower trigger. It follows that the probability that a firm "leaves" a position of relatively high employment in the distribution (via firing) exceeds the probability that the firm returns to it (via hiring). Proposition 3 suggests, then, that fewer firms will occupy positions of relatively high employment in the distributions of the fixed-cost models.

This interpretation is consistent with the differences between fixed-cost models, on one hand, and the frictionless model on the other. But why do the two fixed-cost models (one under foresight, the other under myopia) have different means? And what is the source of differences in dispersion? To answer these questions, we likely have to step outside of the quite general confines of the discussion surrounding Propositions 1-4 and consider the specific form of the stochastic process. An important observation is that productivity is mean-reverting. In this setting, a forward-looking firm that receives a relatively low productivity shock and that chooses to adjust will select a large workforce than the corresponding myopic and frictionless firms. This explains why mean employment is higher under foresight than myopia. It also can account for the difference in dispersion, since it suggests that distribution under foresight will be relatively more compressed.

Lastly, from Table 2, we report the quarterly probability of adjusting. This is very similar across the forward-looking policy (46 percent) and the myopic (45 percent). Thus, the differences in the policy rules do not imply sharply different aggregate adjustment rates. We note briefly that this finding may have implications for empirical work. We suspect, for instance, that an econometrician would be hard-pressed to distinguish between these behavioral policies based on moments of the data, such as adjustment probabilities, that, a priori, would seem to be salient.

Thus far, we have studied aggregate labor demand given some real wage. It is worth remembering that the introduction of general equilibrium always "squeezes" the differences across these models to zero. Figure 6 illustrates this point. It traces out aggregate labor demand schedules for all three models and imposes an upward-sloped supply schedule.[13] Average firm size now varies from 19.85 under a myopic policy to 20.15 in the frictionless model, so the differences are notably smaller.

In the next section, we return to this issue of how general equilibrium forces affect the aggregate dynamics of the models. But as the steady-state moments in Table 2 suggest, differences in the model are rather contained even in partial equilibrium, as predicted by Propositions 1 and 4. This continues to be true out of steady state, as we will see.

3 Aggregate Dynamics

In this section, we show that the symmetry that underlies the approximate aggregate neutrality result of Proposition 4 holds not only in the steady state, but also along the dynamic transition path in the presence of aggregate shocks.

We do this by first characterizing the dynamics of the cross-sectional distribution of employment. A useful feature of the mass balance approach that we used to solve for steady-state outcomes is that it provides a direct description of these dynamics: The change in the

[13] Mulligan (2001) provides a model of heterogeneous reservation wages that implies a smooth, upward-sloped aggregate labor supply curve. As for its elasticity, we assume it is one. This calibration is discussed further in the next section.

mass of firms with employment below some level n between this period and next, $\Delta H(n)'$, is simply equal to the inflow into the mass less the outflow from the mass. Applying this method yields Proposition 5.

Proposition 5 *The law of motion for the density of employment across firms is given by*

$$\begin{aligned}\Delta h(n)' &= \left(1 - \mathcal{H}\left[L^{-1}X(n), X(n)\right] + \mathcal{H}\left[U^{-1}X(n), X(n)\right]\right) h^*(n) \\ &\quad - \left(1 - \mathcal{G}\left[U(n), n\right] + \mathcal{G}\left[L(n), n\right]\right) h(n), \end{aligned} \tag{17}$$

where $h^*(n)$ is the frictionless density of employment, $\mathcal{H}(\nu, \xi) \equiv \Pr[n_{-1} \le \nu | x = \xi]$, and $\mathcal{G}(\xi, \nu) \equiv \Pr[x \le \xi | n_{-1} = \nu]$.

Recalling the discussion preceding equation (14), note that one can restate Proposition 5 as

$$\Delta h(n)' = \Pr(\text{adjust to } n)\, h^*(n) - \Pr(\text{adjust from } n)\, h(n). \tag{18}$$

Intuitively, the inflow into the density of employment at some level n originates from firms that have 1) received an idiosyncratic shock that leads to a desired employment level of n, and 2) inherited an employment level sufficiently different from n such that it is optimal to pay the fixed cost and adjust. Thus, the inflow is equal to the density of firms that would choose an employment level of n absent the adjustment cost, $h^*(n)$, times the probability that they will in fact adjust. Likewise, the outflow from the density of employment at n is simply the share of firms at that level of employment that receives a sufficiently large idiosyncratic shock to adjust away from n. Clearly, in steady state, $\Delta h(n)' = 0$, and equation (17) yields the steady-state outcome in Proposition 3.

Mirroring Proposition 4, it is in turn straightforward to derive a simple approximation to the aggregate dynamics of the model in the case in which the adjustment cost is small, $C \approx 0$. Proposition 6 summarizes the result of this exercise—that the adjustment cost is approximately neutral with respect to the aggregate dynamics of the model.

Proposition 6 *In the presence of a small fixed adjustment cost ($C \approx 0$), the evolution of the distribution of employment across firms is approximated by*

$$\Delta h(n)' \approx -\left[h(n) - h^*(n)\right], \tag{19}$$

which is the frictionless law of motion.

Equation (19) implies that any gap between the distribution of employment and its frictionless counterpart is closed immediately—the aggregate dynamics of the model are approximately jump, and coincide with frictionless dynamics.

The intuition behind Proposition 6 is related to the insights obtained from the steady-state analysis. Suppose a small fixed cost is introduced into an otherwise frictionless environment. At any instant of time, a small fixed cost reduces the outflow of mass from any given employment level n, but also reduces the mass of firms which find it optimal to adjust to that n. For small frictions, these two forces are symmetric and therefore offset each other almost exactly and leave the distribution approximately equal to its frictionless counterpart.

3.1 Numerical Illustration

To explore Proposition 6 quantitatively, we now solve the model numerically out of steady state using the baseline calibration. Throughout, we specify the process of shocks to aggregate productivity p to follow a geometric AR(1)

$$\log p' = \mu_p + \rho_p \log p + \varepsilon_p', \text{ where } \varepsilon_p' \sim N\left(0, \sigma_p^2\right). \tag{20}$$

We calibrate the persistence, ρ_p, and volatility, σ_p, so that the model approximately replicates the persistence and volatility of (de-trended) log aggregate employment. Using time series data on private payroll employment and detrending using the HP filter, we compute an autocorrelation coefficient of 0.96 and a standard deviation of 0.026. Values of $\rho_p = 0.95$ and $\sigma_p = 0.015$ are roughly consistent with these moments (see Table 1). We do this because our goal is not to explain the volatility of aggregate employment. Rather, we seek to compare model outcomes within an environment that is economically relevant. One way of doing that is to generate aggregate outcomes that are comparable to what we observe in the data.

We then compute the impulse response of aggregate employment to an aggregate productivity innovation across three versions of the model: the forward-looking case ($\beta > 0$ and $C > 0$), the myopic case ($\beta = 0$ and $C > 0$), and the frictionless case ($C = 0$).

Solving for the general equilibrium dynamics requires specifying the supply side of the market. We carry out the analysis in this section under two assumptions on labor supply. The first is that the elasticity of labor supply is equal to one. This is the value advocated by Kimball and Shapiro (2010) based on an analysis of survey evidence, and is similar to what is implied by Chang and Kim's (2006) structural model of indivisible labor supply. The second is that labor supply is perfectly elastic, in which case the wage may be treated as fixed (and given by its value in Table 1).[14]

Upward-sloped labor supply We introduce an upward-sloped aggregate labor supply schedule of the simple loglinear form

$$N^s(w) = \psi w^\eta. \tag{21}$$

Equilibrium is thus given by the intersection of demand (15) and the supply schedule, (21).

As noted, we impose that the (Frisch) labor supply elasticity is one, so $\eta = 1$. The intercept, ψ, is set so that mean equilibrium employment remains near 20, consistent with the parameterization in our partial equilibrium analysis above. As shown in Appendix B, this labor supply schedule may be derived from the problem of a "large" household that must allocate its members across market and non-market work. The members suffer different levels of disutility from market work, and the household head selects a cut-off such that all members whose disutility falls below the threshold participate in the labor market.[15] The positive slope in (21) reflects the added disutility borne by the household when it deploys to the labor market another of its members who faces a higher disutility from market work on the margin.[16]

It is well-known that, in the presence of (21), firms face a difficult prediction problem: Because employment is quasi-fixed, firms must forecast the future path of aggregate wages when they decide on the size of their current workforce. Yet the wage depends on aggregate employment, $N \equiv \int nh(n)\,dn$, which in turn depends on the distribution of employment across firms, an infinite-dimensional object.

We adopt Krusell and Smith's (1998) bounded rationality algorithm to solve this problem. Specifically, we assume that firms forecast log aggregate employment using its lag, and the current level of log aggregate productivity,

$$\log N = \theta_0 + \theta_N \log N_{-1} + \theta_p \log p. \tag{22}$$

[14] To clarify: we parameterize the process (20) so that it is consistent with the empirical variation in aggregate log employment, conditional on a labor supply elasticity of one. We do not re-calibrate the model when the elasticity is set to infinity. This allows us to compare impulse responses given a fixed stochastic process for p.

[15] To ease computational burden, we assume that members' utility is linear in consumption. In that case, the marginal utility of wealth is fixed at one (and hence absent from (21)), and there is only one price we have to track. This restriction may be relaxed, though it is not clear to us why it would materially affect the results.

[16] This set-up is similar to that used in Mulligan (2001). See his paper for a more exhaustive treatment of the household's problem.

Using their forecast of aggregate employment implied by equation (22), firms can then forecast future wages.

Figure 7 presents the impulse responses of aggregate employment implied by the forward-looking, myopic and frictionless versions of this dynamic general-equilibrium model. Consistent with the result of Proposition 6, the differences between the impulse responses across the three models are very small—neither the adjustment cost nor forward-looking behavior seem to affect greatly the response of aggregate employment to aggregate shocks. It follows that the forecast equation (22) is very accurate—estimating the equation on model-generated data yields an R^2 in excess of 0.9999. By the same token, the estimated coefficients are very close to what would be expected from the frictionless model, given the calibration of the labor supply elasticity.[17]

Perfectly-elastic labor supply (fixed wages) The second case we examine assumes that the elasticity of labor supply is infinite, so that the wage w is effectively fixed. In doing so, we suppress any general equilibrium adjustment of wages along the transition path. Examining the aggregate dynamics of the model in this environment allows us to isolate features of these dynamics that can be traced to the equilibration of the distribution of employment from those driven by the adjustment of wages in general equilibrium.

With fixed wages, the numerical model is much simpler to solve. In particular, it is no longer necessary to apply the Krusell-Smith bounded rationality algorithm in this context, since the wage no longer needs to be forecasted. Instead, all that needs to be done is to solve for the optimal labor demand policy functions, $L(n,p)$, $X(n,p)$, and $U(n,p)$. A positive innovation to aggregate productivity p shifts these functions downward—for a given level of idiosyncratic productivity, a firm is more likely to hire, less likely to fire, and will select a higher level of employment conditional on adjustment. Thus, the evolution of aggregate productivity p induces shifts in the policy function, which in turn traces out the evolution of the distribution of employment and thereby aggregate employment.

The results of this exercise are illustrated in Figure 7. As in the case with general equilibrium adjustment of wages, the dynamic response of aggregate employment in the presence of adjustment costs lies very close to the frictionless response. Thus, the dynamic neutrality observed in Figure 7 derives in large part from the neutrality of the dynamics of the distribution of employment, as opposed to the adjustment of wages. That said, one can discern small differences between the adjustment paths that were less apparent in the

[17] Specifically, we get $\hat{\theta}_0 = 2.97$, $\hat{\theta}_p = 0.747$, and $\hat{\theta}_N = 0.008$. In the frictionless model, the elasticity with respect to aggregate productivity is 0.735.

case with flexible wages. Intuitively, the upward-sloped labor supply curve attenuates the employment response to changes in aggregate labor demand.

3.2 Discussion

Our finding that the presence of a fixed adjustment cost is approximately neutral with respect to aggregate steady-state outcomes and aggregate dynamics dovetails with the large literature on adjustment costs. While we have focused on the adjustment of employment, the model has direct analogues in the literatures on investment dynamics and price rigidity, so we will touch on those literatures also in the course of our discussion.

Within the literature on dynamic labor demand, there are two papers in particular that are related to ours. The first is King and Thomas (2006). Their work differs from ours principally in that they do not model productive heterogeneity. Rather, they assume that firms face stochastic fixed costs of adjusting: in each period, firms take an i.i.d. draw from a distribution of fixed costs. If there is exogenous attrition, then firms will choose to replenish their stock of workers if the fixed cost drawn that period is sufficiently small.

King and Thomas find that the response of aggregate employment displays a slight hump shape, which is reminiscent of a partial-adjustment model. This is true both in partial and general equilibrium. It is also true despite the fact that the typical fixed cost is small: we compute that the mean adjustment cost in their calibration is just over one percent of revenue. Why, then, do King and Thomas find different dynamics than reported in section 3.1?

We suspect the reason lies in the absence of productive heterogeneity. In their model, the distribution of employment is highly left-skewed. There is a mass of firms which have just adjusted and so have the desired number of workers; all other firms lie to the left of this point in the distribution. They have seen their workforces partially depleted by attrition and now have fewer than the ideal number of workers. The frictionless counterpart to this model has only aggregate productivity variation, so the distribution is degenerate. Hence, the invariance result of Proposition 4 simply does not apply.

The introduction of idiosyncratic risk induces a less skewed distribution of employment. Put another way, when there is productive heterogeneity, there is a probability that a firm's future employment will be too low *and* a probability that it will be too high relative to its next-period productivity realization). In this more symmetric setting, our myopic approximation will be more accurate; it relies, after all, on the approximate symmetry in the marginal value of employment around the reset function. The myopic approximation is, in turn, a pathway to the aggregate near-invariance results.

A second paper of interest is Bachmann (2009), who also finds that a fixed adjustment cost model induces sluggish, partial-adjustment-like dynamics in aggregate employment. Bachmann does allow for productive heterogeneity, and his calibration of idiosyncratic risk is comparable to that used in this paper. What is different in Bachmann's calibration, though, is that the share of firms which adjust per quarter is quite low, at 6.4 percent. Thus, the contrast between our results and those in Bachmann serves to highlight the circumstances in which neutrality can be expected to arise, namely when the adjustment cost is sufficiently small relative to the volatility of idiosyncratic shocks. This indicates that it will be fruitful to focus empirical work on obtaining robust estimates of these two critical parameters in order to infer the role of lumpy microeconomic adjustment in aggregate employment dynamics.

4 Conclusion

This paper has analyzed labor demand in the presence of a fixed cost of adjustment. It has shown that for small, though plausible, fixed costs, the optimal labor demand rule is well approximated by the myopic policy. The key to this result was a certain symmetry embedded in the firm's problem: When the adjustment cost is small, the positive effect of carrying a marginally higher level of employment into a state of the world where productivity is higher is just offset by the negative effect of carrying an excess work into a state of the world where productivity is lower. In expectation, the marginal effect of current employment decisions on future profits is thus approximately zero.

We then used this myopic approximation as a vehicle through which to investigate the aggregate consequences of a fixed cost. We found that, for a sufficiently small friction, aggregate employment approximately tracks its frictionless counterpart. The intuition for this result is again bound up with a form of symmetry: the presence of an adjustment cost both reduces the probability of a firm adjusting *to* an employment level and the probability of that firm adjusting away *from* that level. In the neighborhood of a small adjustment cost, these two opposing effects cancel, implying the mass of firms occupying each position in the employment distribution is approximately unchanged from what is observed in the frictionless model. These approximations were shown to hold up reasonably well in a calibrated model.

Under what circumstances, then, do fixed costs matter for aggregate outcomes? The analytical results of this paper point, of course, to the importance of the magnitude of the friction. This indicates that perhaps more work ought to be done, along the lines of Cooper, Haltiwanger, and Willis (2007) and Bloom (2009), to identify the sizes of various employment adjustment frictions. For instance, fixed and linear costs of adjustment both induce inaction and may be able to broadly account for the basic shape of the cross-sectional employment growth distribution. But these costs may have very different implications for aggregate dynamics. Elsby and Michaels (2010), for instance, find that aggregate employment in a model with a kinked (proportional) adjustment cost exhibits some persistence, in contrast to the approximately jump-dynamics in the fixed-cost model.[18]

Another key structural parameter is the variance of idiosyncratic productivity. As King and Thomas (2006) show in the limiting case where this variance is zero, it is possible to generate partial adjustment-like dynamics within a fixed-cost model. This indicates that our analytical results may not provide as good of a guide if this idiosyncratic variance is in fact rather small. Studies to date seem to suggest the opposite but the importance of this parameter for *aggregate* dynamics merits continued study.

Lastly, Caballero, Engel, and Haltiwanger (1997) emphasized that the nonlinearities of a fixed-cost model may not show through in aggregate data unless the aggregate impulse is very large. In additional numerical work, we did not find this to be the case for small frictions. But if fixed costs are in fact large, they may interact with big aggregate shocks to generate interesting dynamics.

5 References

Bachmann, Ruediger, "Understanding Jobless Recoveries," mimeo, University of Michigan, 2009.

Bloom, Nicholas, "The Impact of Uncertainty Shocks," *Econometrica*, Vol. 77, No. 3 (May 2009), pp. 623-685.

Caballero, R., E. Engel, and J. Haltiwanger, "Aggregate Employment Dynamics: Building from Microeconomic Evidence," *American Economic Review*, Vol. 87, No. 1 (March 1997), pp. 115–37.

Caballero, R. and E. Engel, "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach," *Econometrica*, Vol. 67, No.4 (July 1999), pp. 783-826.

Caballero, R. and E. Engel, "Three Strikes and You're Out: Reply to Cooper and Willis," *American Economic Review*, Vol. 94, No.4 (Sept. 2004), pp. 1238-1244.

[18]Elsby and Michaels (2010) consider an extension of Mortensen and Pissarides (1994) to the case where production displays decreasing returns to scale. The cost of vacancy posting in this setting represents a proportional cost of adjusting employment.

Chang, Yongsung and Sun-Bin Kim, "From Individual to Aggregate Labor Supply: A Quantitative Analysis based on a Heterogeneous-Agent Macroeconomy" *International Economic Review*, Vol. 47, No. 1 (2006), pp. 1-27.

Cooper, Russell, John Haltiwanger, and Jonathan L. Willis, "Dynamics of Labor Demand: Evidence from Plant-level Observations and Aggregate Implications," NBER Working Paper No. 10297, 2004.

Cooper, Russell, and Jonathan L. Willis, "A Comment on the Economics of Labor Adjustment: Mind the Gap," *American Economic Review*, Vol. 94, No.4 (Sept. 2004), pp. 1223-1237.

Cooper, Russell, John Haltiwanger, and Jonathan L. Willis, "Implications of Search Frictions: Matching Aggregate and Establishment-level Observations," Working Paper No. 13115, National Bureau of Economic Research, 2007.

Davis, Steven J., and John Haltiwanger, "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, Vol. 107, No. 3 (Aug. 1992), pp. 819-864.

Foster, Lucia, John Haltiwanger, and Chad Syverson, "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review*, Vol. 98, No. 1 (March 2008), pp. 394-425.

Gertler, Mark and John Leahy, "A Phillips Curve with an Ss Foundation," *Journal of Political Economy*, Vol. 116, No. 3 (June 2008), pp. 533-572.

Hamermesh, Daniel, "Labor Demand and the Structure of Adjustment Costs," *American Economic Review*, Vol. 79, No.4, (Sept. 1989), pp. 674-689.

Judd, Kenneth, *Numerical Methods in Economics*, Cambridge, MA: The MIT Press, 1998.

Kahn, Aubhik and Julia Thomas, "Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics," *Econometrica*, Vol. 76, No.2 (March 2008), pp. 395-436.

Kimball, Miles and Matthew Shapiro, "Labor Supply: Are Income and Substitution Effects Both Large or Both Small?" mimeo, University of Michigan, 2010.

King, Robert and Julia Thomas, "Partial Adjustment Without Apology" *International Economic Review,* Vol. 47, No. 3 (August 2006), pp. 779-809.

Krusell, Per, and Anthony Smith, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106:5, 867-896.

6 Appendix

A Proofs

Lemma 1 *In the presence of a small lump-sum fixed adjustment cost, the adjustment triggers and their inverses are approximately equal to*

$$\begin{aligned} L(n) &\approx X(n) - \gamma(n)\sqrt{C}, \qquad U(n) \approx X(n) + \gamma(n)\sqrt{C}, \text{ and} \qquad (23) \\ L^{-1}(x) &\approx X^{-1}(x) + \bar{\gamma}(x)\sqrt{C},\ U^{-1}(n) \approx X^{-1}(n) - \bar{\gamma}(n)\sqrt{C}, \qquad (24) \end{aligned}$$

where $\gamma(\tilde{n}) \equiv \sqrt{2/\Delta_{xx}(\tilde{n}, X(\tilde{n}))}$, *and* $\bar{\gamma}(\tilde{n}) \equiv \sqrt{2/\Delta_{11}(X^{-1}(x), x)}$.

Proof of Lemma 1. Recall that the adjustment triggers satisfy the value matching condition, $\Delta(\tilde{n}, x') \equiv \Pi^{\Delta}(x') - \Pi^{0}(\tilde{n}, x') = C$. In the presence of $C \approx 0$, we may restrict our focus to a second-order approximation to $\Delta(\tilde{n}, x')$ around $x' = X(\tilde{n})$:

$$\Delta(\tilde{n}, x') \approx \Delta(\tilde{n}, X(\tilde{n})) + \Delta_x(\tilde{n}, X(\tilde{n}))(x' - X(\tilde{n})) + \frac{1}{2}\Delta_{xx}(\tilde{n}, X(\tilde{n}))(x' - X(\tilde{n}))^2. \quad (25)$$

The first and second terms on the right side are zero by optimality. Setting $\Delta(\tilde{n}, x') = C$, it follows that the triggers are as stated.

The inverse triggers may be derived symmetrically by approximating $\Delta(\tilde{n}, x')$ around $\tilde{n} = X^{-1}(x)$:

$$\Delta(\tilde{n}, x') \approx \Delta\left(X^{-1}(x), x\right) + \Delta_1\left(X^{-1}(x), x\right)\left(\tilde{n} - X^{-1}(x)\right) + \frac{1}{2}\Delta_{11}\left(X^{-1}(x), x\right)\left(\tilde{n} - X^{-1}(x)\right)^2. \quad (26)$$

Again, optimality implies the first two terms in the expansion are zero. Setting $\Delta(\tilde{n}, x') = C$ yields the stated inverse triggers. ■

Proof of Proposition 1. We seek to show that the expected future value of the firm, $\mathcal{F}(\tilde{n}, x) \equiv \int \Pi(\tilde{n}, x')\, dG(x'|x)$, is approximately independent of n for sufficiently small C. To begin, we partition the forward value into parts associated with each of the three continuation regimes—firing, inaction, and hiring:

$$\begin{aligned} \mathcal{F}(\tilde{n}, x) = & \int_0^{L(\tilde{n})} \left[\Pi^{\Delta}(x') - C\right] dG(x'|x) + \int_{L(\tilde{n})}^{U(\tilde{n})} \Pi^0(\tilde{n}, x')\, dG(x'|x) \\ & + \int_{U(\tilde{n})}^{\infty} \left[\Pi^{\Delta}(x') - C\right] dG(x'|x). \qquad (27) \end{aligned}$$

Differentiating, and using the value-matching conditions in (6) to eliminate the derivatives of the limits of integration yields:

$$\mathcal{F}_1(\tilde{n}, x) = \int_{L(\tilde{n})}^{U(\tilde{n})} \Pi_1^0(\tilde{n}, x')\, dG(x'|x). \quad (28)$$

Consider a first-order approximation to the latter, $\mathcal{F}_1(\tilde{n}, x) \approx \mathcal{F}_1(\tilde{n}, x)|_{C=0} + \mathcal{F}_{1C}(\tilde{n}, x)|_{C=0} \cdot C$. The leading term is zero—in the absence of an adjustment friction, the problem is static.

The derivative, $\mathcal{F}_{1C}(\tilde{n}, x)$, is more difficult to characterize. Our approach is to develop a recursion for this and show that $\mathcal{F}_{1C}(\tilde{n}, x)|_{C=0} = 0$ is implied by the recursion. To begin, differentiate (28) with respect to C, and note from equation (3) that $\Pi^0_{1C}(\tilde{n}, x') = \beta(1-\delta)\mathcal{F}_{1C}((1-\delta)\tilde{n}, x')$ yields

$$\begin{aligned}\mathcal{F}_{1C}(\tilde{n}, x) &= \beta(1-\delta)\int_{L(\tilde{n})}^{U(\tilde{n})} \mathcal{F}_{1C}((1-\delta)\tilde{n}, x')\, dG(x'|x) \\ &\quad + \Pi^0_1(\tilde{n}, U(\tilde{n}))\, g(U(\tilde{n})|x)\frac{\partial U(\tilde{n})}{\partial C} \\ &\quad - \Pi^0_1(\tilde{n}, L(\tilde{n}))\, g(L(\tilde{n})|x)\frac{\partial L(\tilde{n})}{\partial C}. \qquad (29)\end{aligned}$$

We conjecture that $\mathcal{F}_{1C}(\tilde{n}, x)|_{C=0} = 0$ and show that the recursion, (29), verifies this when evaluated at $C = 0$. Under the conjecture, equation (29) implies $\mathcal{F}_1(\tilde{n}, x) \approx 0$. Thus, (3) implies $\Pi^0_1(\tilde{n}, x') \approx px'F_n(\tilde{n}) - w$. Moreover, if $\mathcal{F}_1(\tilde{n}, x) \approx 0$, then the first-order condition (5) for an adjusting firm yields the simple labor demand rule, $pxF_n(n) = w$. Therefore, the policy function is given by $X(n) = w/[pF_n(n)]$. Substitution then implies

$$\begin{aligned}\mathcal{F}_{1C}(\tilde{n}, x) &\approx \beta(1-\delta)\int_{L(\tilde{n})}^{U(\tilde{n})} \mathcal{F}_{1C}((1-\delta)\tilde{n}, x')\, dG(x'|x) \\ &\quad + pF_n(\tilde{n})[U(\tilde{n}) - X(\tilde{n})]\, g(U(\tilde{n})|x)\frac{\partial U(\tilde{n})}{\partial C} \\ &\quad - pF_n(\tilde{n})[L(\tilde{n}) - X(\tilde{n})]\, g(L(\tilde{n})|x)\frac{\partial L(\tilde{n})}{\partial C}. \qquad (30)\end{aligned}$$

At this point, we turn our attention to the triggers, $L(\tilde{n})$ and $U(\tilde{n})$. Using equation (23) from Lemma 1, and their derivatives, we can write[19]

$$\begin{aligned}\mathcal{F}_{1C}(\tilde{n}, x) &\approx \beta(1-\delta)\int_{L(\tilde{n})}^{U(\tilde{n})} \mathcal{F}_{1C}((1-\delta)\tilde{n}, x')\, dG(x'|x) \qquad (31) \\ &\quad + \frac{1}{2}pF_n(\tilde{n})\gamma(\tilde{n})^2\left[g\left(X(\tilde{n}) + \gamma(\tilde{n})\sqrt{C}|x\right) - g\left(X(\tilde{n}) - \gamma(\tilde{n})\sqrt{C}|x\right)\right].\end{aligned}$$

As the adjustment friction approaches zero, $C \to 0$, the triggers converge. Thus, given the conjecture, the integral on the right side converges to zero. In addition, as long as the p.d.f. $g(\cdot)$ is continuous, the bottom line also converges to zero. We conclude that the conjecture is confirmed: $\mathcal{F}_{1C}(\tilde{n}, x)|_{C=0} = 0$. ■

Proof of Corollary 1. The introduction of a size-dependent adjustment cost does not substantively affect the derivation of the recursion in the cross-partial, $\mathcal{F}_{1C}(\tilde{n}, x)$: One may

[19] When we take these derivatives, recall that, since $X(n) = w/[pF_n(n)]$ under the conjecture, it is therefore independent of C.

replace $\mathcal{F}_{1C}(\tilde{n}, x)$ with $\mathcal{F}_{1c}(\tilde{n}, x)$, and this portion of the proof holds *mutatis mutandis*. Only the analysis of the triggers must be revisited. The value matching condition instead becomes $\cdot(\tilde{n}, x) = c(x')^{\frac{1}{1-\alpha}}$. Thus, the triggers are given implicitly by

$$L(\tilde{n}) \approx X(\tilde{n}) - \gamma(\tilde{n})\sqrt{c}L(\tilde{n})^{\frac{1}{2}\frac{1}{1-\alpha}}, \text{ and } U(\tilde{n}) \approx X(\tilde{n}) + \gamma(\tilde{n})\sqrt{c}U(\tilde{n})^{\frac{1}{2}\frac{1}{1-\alpha}}. \tag{32}$$

Differentiating with respect to c yields

$$\frac{\partial L(\tilde{n})}{\partial c} \approx -\frac{\gamma(\tilde{n})L(\tilde{n})^{\frac{1}{2}\frac{1}{1-\alpha}}}{2+\gamma(\tilde{n})\frac{1}{1-\alpha}\sqrt{c}L(\tilde{n})^{\frac{1}{2}\frac{2\alpha-1}{1-\alpha}}}\frac{1}{\sqrt{c}}, \text{ and } \frac{\partial U(\tilde{n})}{\partial c} \approx \frac{\gamma(\tilde{n})U(\tilde{n})^{\frac{1}{2}\frac{1}{1-\alpha}}}{2-\gamma(\tilde{n})\frac{1}{1-\alpha}\sqrt{c}U(\tilde{n})^{\frac{1}{2}\frac{2\alpha-1}{1-\alpha}}}\frac{1}{\sqrt{c}}. \tag{33}$$

Substituting the triggers and their derivatives into the recursion (30), we obtain

$$\begin{aligned}\mathcal{F}_{1c}(\tilde{n}, x) \approx\ & \beta(1-\delta)\int_{L(\tilde{n})}^{U(\tilde{n})}\mathcal{F}_{1c}((1-\delta)\tilde{n}, x')\,dG(x'|x) \qquad (34)\\ & +pF_n(\tilde{n})\gamma(\tilde{n})^2\left[\begin{array}{l} g\left(X(\tilde{n})+\gamma(\tilde{n})\sqrt{c}U(\tilde{n})^{\frac{1}{2}\frac{1}{1-\alpha}}|x\right)\frac{U(\tilde{n})^{\frac{1}{1-\alpha}}}{2-\gamma(\tilde{n})\frac{1}{1-\alpha}\sqrt{c}U(\tilde{n})^{\frac{1}{2}\frac{2\alpha-1}{1-\alpha}}}\\ -g\left(X(\tilde{n})-\gamma(\tilde{n})\sqrt{c}L(\tilde{n})^{\frac{1}{2}\frac{1}{1-\alpha}}|x\right)\frac{L(\tilde{n})^{\frac{1}{1-\alpha}}}{2+\gamma(\tilde{n})\frac{1}{1-\alpha}\sqrt{c}L(\tilde{n})^{\frac{1}{2}\frac{2\alpha-1}{1-\alpha}}}\end{array}\right].\end{aligned}$$

Evaluating this at $c = 0$, the integral again collapses to zero, and the term in brackets becomes $\frac{1}{2}X(\tilde{n})^{\frac{1}{1-\alpha}}[g(X(\tilde{n})|x) - g(X(\tilde{n})|x)] = 0$, where we have used (32) evaluated at $c = 0$. This confirms that Proposition 1 also holds when the fixed cost takes the form $C = cx^{\frac{1}{1-\alpha}}$. ■

Proof of Proposition 2. The derivation of $X(n)$ follows directly from the static first-order condition. The adjustment triggers are the values of idiosyncratic productivity x that satisfy the value-matching condition,

$$pxn^{*\alpha} - wn^* = pxn^\alpha - wn + cR(x), \tag{35}$$

where $n^* = (\alpha px/w)^{1/(1-\alpha)}$ is optimal employment (conditional on adjusting). Frictionless revenue therefore is equal to $pxn^{*\alpha} = (px)^{1/(1-\alpha)}(\alpha/w)^{\alpha/(1-\alpha)}$, and so the adjustment cost $C(x) = c(px)^{1/(1-\alpha)}(\alpha/w)^{\alpha/(1-\alpha)}$. After substitution, the value-matching condition becomes

$$(1-\alpha-c)(px)^{\frac{1}{1-\alpha}}(\alpha/w)^{\frac{\alpha}{1-\alpha}} = pxn^\alpha - wn. \tag{36}$$

Conjecture that the roots of the latter equation are of the form $x = \mathcal{Z}n^{1-\alpha}$. Imposing the conjecture, and noting that $\mathcal{X} \equiv w/\alpha p$, yields the stated result. ■

Proof of Proposition 3. The following summarizes the proof for the case where there is no exogenous workforce attrition, $\delta = 0$. The general proof is in progress, but follows a very similar strategy.

We wish to derive the flows in and out of the mass of firms with employment below some number m, $H(m)$. Consider first the inflow into that mass—i.e. the mass of firms that cuts employment from above m to below m. The latter is given by

$$\begin{aligned}\text{Inflow into } H(m) &= \int_{L^{-1}X(m)}^{\infty} \mathcal{G}\left[X(m), n_{-1}\right] dH(n_{-1}) \\ &\quad + \int_{m}^{L^{-1}X(m)} \mathcal{G}\left[L(n_{-1}), n_{-1}\right] dH(n_{-1}), \end{aligned} \tag{37}$$

where $\mathcal{G}(\xi, \nu) \equiv \Pr\left[x \leq \xi | n_{-1} = \nu\right]$. To understand this, first fix a level of lagged employment, n_{-1}. There are two sets of inflows corresponding to the following two cases:

1) If $m < X^{-1}L(n_{-1})$, so that $n_{-1} > L^{-1}X(m)$, it can be seen from the optimal policy in Figure 1 that the probability of reducing employment below m will be $\mathcal{G}\left[X(m), n_{-1}\right]$.

2) If $m \in \left[X^{-1}L(n_{-1}), n_{-1}\right]$, so that $n_{-1} \in \left[m, L^{-1}X(m)\right]$, Figure 1 suggests that the probability of reducing employment below m will be $\mathcal{G}\left[L(n_{-1}), n_{-1}\right]$.

The first and second terms in (37) respectively aggregate these inflows over the relevant values of n_{-1}. Following a similar logic, the outflow from $H(m)$ can be expressed as

$$\begin{aligned}\text{Outflow from } H(m) &= \int_{U^{-1}X(m)}^{m} \left(1 - \mathcal{G}\left[U(n_{-1}), n_{-1}\right]\right) dH(n_{-1}) \\ &\quad + \int_{0}^{U^{-1}X(m)} \left(1 - \mathcal{G}\left[X(m), n_{-1}\right]\right) dH(n_{-1}). \end{aligned} \tag{38}$$

Setting the inflow into the mass equal to the outflow from that mass, noting that $G\left[X(m)\right] = \int_0^{\infty} \mathcal{G}\left[X(m), n_{-1}\right] dH(n_{-1})$, and solving for $H(m)$ yields

$$\begin{aligned} H(m) &= G\left[X(m)\right] - \int_{U^{-1}X(m)}^{L^{-1}X(m)} \mathcal{G}\left[X(m), n_{-1}\right] dH(n_{-1}) \qquad (39) \\ &\quad + \int_{m}^{L^{-1}X(m)} \mathcal{G}\left[L(n_{-1}), n_{-1}\right] dH(n_{-1}) + \int_{U^{-1}X(m)}^{m} \mathcal{G}\left[U(n_{-1}), n_{-1}\right] dH(n_{-1}). \end{aligned}$$

Differentiating with respect to m and cancelling terms implies that the density of employment across firms is given by

$$h(m) = h^*(m) \frac{1 - \int_{U^{-1}X(m)}^{L^{-1}X(m)} \frac{\mathcal{G}_1[X(m), n_{-1}]}{g[X(m)]} dH(n_{-1})}{1 - \mathcal{G}\left[U(m), m\right] + \mathcal{G}\left[L(m), m\right]}, \tag{40}$$

where $h^*(m) \equiv g[X(m)]X'(m)$ is the frictionless density of employment. To obtain the stated expression stated, note that one can use Bayes' rule to write the distribution of lagged employment conditional on current productivity as

$$\mathcal{H}(\nu, \xi) \equiv \Pr\left[n_{-1} \leq \nu | x = \xi\right] = \int_0^{\nu} \frac{\mathcal{G}_1(\xi, \tilde{\nu})}{g(\xi)} dH(\tilde{\nu}), \tag{41}$$

so that $\mathcal{H}\left[L^{-1}X\left(m\right),X\left(m\right)\right]-\mathcal{H}\left[U^{-1}X\left(m\right),X\left(m\right)\right]=\int_{U^{-1}X(m)}^{L^{-1}X(m)}\frac{\mathcal{G}_1[X(m),n_{-1}]}{g[X(m)]}dH\left(n_{-1}\right)$. ■

Lemma 2 *The density of idiosyncratic productivity conditional on lagged employment satisfies the recursion*

$$\mathcal{G}_1\left(\xi,\nu\right)=\pi^0\frac{\int_{L(\nu)}^{U(\nu)}g\left(\xi|x_{-1}\right)\mathcal{G}_1\left(x_{-1},\nu\right)dx_{-1}}{\mathcal{G}\left[U\left(\nu\right),\nu\right]-\mathcal{G}\left[L\left(\nu\right),\nu\right]}+\left(1-\pi^0\right)g\left[\xi|X\left(\nu\right)\right]. \tag{42}$$

It follows that $\mathcal{G}_1\left(x,n_{-1}\right)$ *is bounded and continuous in* x.

Proof of Lemma 2. First note that we may write $\mathcal{G}\left(\xi,\nu\right)=\int G\left(\xi|x_{-1}\right)d\mathbb{G}\left(x_{-1}|\nu\right)$, where

$$\begin{aligned}\mathbb{G}\left(\tilde{\xi}|\nu\right) &\equiv \Pr\left[x_{-1}\leq\tilde{\xi}|n_{-1}=\nu\right]\\ &= \Pr\left[x_{-1}\leq\tilde{\xi}|n_{-1}=\nu=n_{-2}\right]\Pr\left[n_{-1}=n_{-2}\right]\\ &\quad+\Pr\left[x_{-1}\leq\tilde{\xi}|n_{-1}=\nu,n_{-1}\neq n_{-2}\right]\Pr\left[n_{-1}\neq n_{-2}\right]. \end{aligned}\tag{43}$$

In the event that the firm adjusted last period, $n_{-1}\neq n_{-2}$, we know that the firm would have adjusted so that $x_{-1}=X\left(n_{-1}\right)$. Thus,

$$\Pr\left[x_{-1}\leq\tilde{\xi}|n_{-1}=\nu,n_{-1}\neq n_{-2}\right]=\mathbf{1}\left[\tilde{\xi}\geq X\left(\nu\right)\right]. \tag{44}$$

In the case in which the firm did not adjust last period, we know that $n_{-2}=\nu$. That information alone implies that x_{-1} will be distributed according to the c.d.f. of $x_{-1}|n_{-2}$, that is $\mathcal{G}$. In addition, however, we also know that $n_{-1}=\nu$. This implies that $x_{-1}\in\left[L\left(v\right),U\left(\nu\right)\right]$, but is otherwise uninformative on the distribution of x_{-1}. Thus,

$$\Pr\left[x_{-1}\leq\tilde{\xi}|n_{-1}=\nu=n_{-2}\right]=\frac{\mathcal{G}\left(\tilde{\xi},\nu\right)-\mathcal{G}\left[L\left(\nu\right),\nu\right]}{\mathcal{G}\left[U\left(\nu\right),\nu\right]-\mathcal{G}\left[L\left(\nu\right),\nu\right]}. \tag{45}$$

Defining $\pi^0\equiv\Pr\left[n_{-1}=n_{-2}\right]=\int\left(\mathcal{G}\left[U\left(\tilde{\nu}\right),\tilde{\nu}\right]-\mathcal{G}\left[L\left(\tilde{\nu}\right),\tilde{\nu}\right]\right)dH\left(\tilde{\nu}\right)$, we can therefore write

$$\mathbb{G}\left(\tilde{\xi}|\nu\right)=\pi^0\frac{\mathcal{G}\left(\tilde{\xi},\nu\right)-\mathcal{G}\left[L\left(\nu\right),\nu\right]}{\mathcal{G}\left[U\left(\nu\right),\nu\right]-\mathcal{G}\left[L\left(\nu\right),\nu\right]}+\left(1-\pi^0\right)\mathbf{1}\left[\tilde{\xi}\geq X\left(\nu\right)\right]. \tag{46}$$

Substituting into the definition of $\mathcal{G}\left(\xi,\nu\right)$ yields the stated result. Boundedness and continuity in ξ follow from the boundedness and continuity of the exogenous density of idiosyncratic shocks $g\left(x'|x\right)$. ■

Lemma 3 *Consider integrals of the form,* $\lim_{c\to0}\frac{1}{c}\int_a^{a+bc}f\left(x;c\right)dx$. *For any given* $c\in(0,\infty)$, *pick* b *so that* $f\left(x;c\right)$ *is bounded on* $\left[a,a+bc\right]$ *with* $\sup_{x\in[a,a+bc]}f\left(x;c\right)=\overline{M}\left(c\right)$

and $\inf_{x\in[a,a+bc]} f(x;c) = \underline{M}(c)$. *Then if* $\lim_{c\to 0}\underline{M}(c) = \lim_{c\to 0}\overline{M}(c) = M^*$, *it follows that*

$$\lim_{c\to 0}\frac{1}{c}\int_a^{a+bc} f(x;c)\,dx = bM^*.$$

Proof of Lemma 3. By the Comparison Theorem (William Wade, 5.2), $\int_a^{a+bc} f(x;c)\,dx \leq bc\overline{M}(c)$. It follows that $\frac{1}{c}\int_a^{a+bc} f(x;c)\,dx \leq b\overline{M}(c)$ and $\lim_{c\to 0}\frac{1}{c}\int_a^{a+bc} f(x;c)\,dx \leq b\lim_{c\to 0}\overline{M}(c)$. Conversely, if $\inf_{x\in[a,a+bc]} f(x;c) = \underline{M}(c)$, then $b\lim_{c\to 0}\underline{M}(c) \leq \lim_{c\to 0}\frac{1}{c}\int_a^{a+bc} f(x;c)\,dx$. Moreover, by the Squeeze Theorem, if $\lim_{c\to 0}\underline{M}(c) = \lim_{c\to 0}\overline{M}(c) = M^*$, then the statement in the Lemma holds. ■

Proof of Proposition 4. Consider a first-order Taylor series approximation to $H(n)$ around $C=0$,

$$H(n) \approx H^*(n) + \left.\frac{\partial H(n)}{\partial C}\right|_{C=0}\cdot C, \tag{47}$$

where $H^*(n)$ is the frictionless distribution of employment. We want to evaluate $\partial H(n)/\partial C|_{C=0}$. To do this, we conjecture that $\partial H(n)/\partial C|_{C=0} = 0 = \partial\mathcal{G}(x,n_{-1})/\partial C|_{C=0}$, and then verify that this is implied by the recursions for $H(n)$ and $\mathcal{G}(x,n_{-1})$. Note also that, in the myopic case, $\partial X(n)/\partial C = 0$.

We first verify that $\partial\mathcal{G}/\partial C|_{C=0} = 0$. Using the conjecture, and Lemmas 1-3, we can write

$$\begin{aligned}
\frac{\partial\mathcal{G}(\xi,\nu)}{\partial C} &\approx \frac{\partial\pi^0}{\partial C}\frac{\int_{X(\nu)-\gamma(\nu)\sqrt{C}}^{X(\nu)+\gamma(\nu)\sqrt{C}}\left(G(\xi|x_{-1}) - G[\xi|X(\nu)]\right)\mathcal{G}_1(x_{-1},\nu)\,dx_{-1}}{\mathcal{G}[U(\nu),\nu] - \mathcal{G}[L(\nu),\nu]} \\
&\quad + \pi^0\frac{\gamma(\nu)}{2\sqrt{C}}\left(G[\xi|U(\nu)]\,\mathcal{G}_1[U(\nu),\nu] + G[\xi|L(\nu)]\,\mathcal{G}_1[L(\nu),\nu]\right) \\
&\equiv A + B.
\end{aligned} \tag{48}$$

We then proceed term by term. In the neighborhood of $C=0$, we can write

$$\begin{aligned}
\pi^0 &\approx \int\left(\mathcal{G}\left[X(\tilde{\nu}) + \gamma(\tilde{\nu})\sqrt{C},\tilde{\nu}\right] - \mathcal{G}\left[X(\tilde{\nu}) - \gamma(\tilde{\nu})\sqrt{C},\tilde{\nu}\right]\right)dH(\tilde{\nu}) \\
&= \int\left[\begin{array}{c}\mathcal{G}[X(\tilde{\nu}),\tilde{\nu}] + \mathcal{G}_1[X(\tilde{\nu}),\tilde{\nu}]\,\gamma(\tilde{\nu})\sqrt{C} + \frac{1}{2}\mathcal{G}_{11}[X(\tilde{\nu}),\tilde{\nu}]\,\gamma(\tilde{\nu})^2 C + ... \\ -\mathcal{G}[X(\tilde{\nu}),\tilde{\nu}] + \mathcal{G}_1[X(\tilde{\nu}),\tilde{\nu}]\,\gamma(\tilde{\nu})\sqrt{C} - \frac{1}{2}\mathcal{G}_{11}[X(\tilde{\nu}),\tilde{\nu}]\,\gamma(\tilde{\nu})^2 C + ...\end{array}\right]dH(\tilde{\nu}) \\
&= \int\left(2\mathcal{G}_1[X(\tilde{\nu}),\tilde{\nu}]\,\gamma(\tilde{\nu})\sqrt{C} + ...\right)dH(\tilde{\nu}) \\
&= \mathcal{O}\left(C^{1/2}\right).
\end{aligned} \tag{49}$$

It follows that $\partial\pi^0/\partial C = \mathcal{O}\left(C^{-1/2}\right)$, and that $\mathcal{G}[U(\tilde{\nu}),\tilde{\nu}] - \mathcal{G}[L(\tilde{\nu}),\tilde{\nu}] = \mathcal{O}\left(C^{1/2}\right)$. Following an analogous method, after some tedious alegebra one can show

$$\begin{aligned}\int_{X(\nu)-\gamma(\nu)\sqrt{C}}^{X(\nu)+\gamma(\nu)\sqrt{C}} \left(G(\xi|x_{-1}) - G[\xi|X(\nu)]\right)\mathcal{G}_1(x_{-1},\nu)\,dx_{-1} &= \mathcal{O}\left(C^{3/2}\right), \text{ and} \\ G[\xi|U(\nu)]\,\mathcal{G}_1[U(\nu),\nu] + G[\xi|L(\nu)]\,\mathcal{G}_1[L(\nu),\nu] &= \mathcal{O}\left(C^{1/2}\right). \end{aligned} \tag{50}$$

To summarize, it follows that $A = \mathcal{O}\left(C^{-1/2}\right)\mathcal{O}\left(C^{3/2}\right)/\mathcal{O}\left(C^{1/2}\right) = \mathcal{O}\left(C^{1/2}\right)$, and $B = \mathcal{O}\left(C^{1/2}\right)\mathcal{O}\left(C^{-1/2}\right)\mathcal{O}\left(C^{1/2}\right) = \mathcal{O}\left(C^{1/2}\right)$. Thus, $\partial\mathcal{G}(\xi,\nu)/\partial C = \mathcal{O}\left(C^{1/2}\right) \to 0$ as $C \to 0$, as conjectured.

We now turn to verifying that $\partial\mathcal{H}/\partial C|_{C=0} = 0$. Noting that $\partial\mathcal{G}/\partial C|_{C=0} = 0$, in the neighborhood of $C = 0$ we can use (39) to write

$$\begin{aligned}\frac{\partial H(n)}{\partial C} &\approx \int_{n}^{L^{-1}X(n)} \mathcal{G}_1[L(n_{-1}),n_{-1}]\,\frac{\partial L(n_{-1})}{\partial C}dH(n_{-1}) \\ &\quad + \int_{U^{-1}X(n)}^{n} \mathcal{G}_1[U(n_{-1}),n_{-1}]\,\frac{\partial U(n_{-1})}{\partial C}dH(n_{-1}) \\ &\equiv \mathrm{I} + \mathrm{II}\end{aligned} \tag{51}$$

Using Lemmas 1, 2 and 3, we can write the limit of the first term on the right-hand side of (51) as

$$\begin{aligned}\lim_{C\to 0}\mathrm{I} &= -\frac{1}{2}\lim_{C\to 0}\frac{1}{\sqrt{C}}\int_{n}^{n+\bar{\gamma}(X(n))\sqrt{C}} \mathcal{G}_1\left[X(n_{-1}) - \gamma(n_{-1})\sqrt{C}, n_{-1}\right]\gamma(n_{-1})\,dH(n_{-1}) \\ &= -\frac{1}{2}\bar{\gamma}(X(n))\,\mathcal{G}_1[X(n),n]\,\gamma(n).\end{aligned} \tag{52}$$

Analogously, we can write

$$\begin{aligned}\lim_{C\to 0}\mathrm{II} &= \frac{1}{2}\lim_{C\to 0}\frac{1}{\sqrt{C}}\int_{n-\bar{\gamma}(X(n))\sqrt{C}}^{n} \mathcal{G}_1\left[X(n_{-1}) + \gamma(n_{-1})\sqrt{C}, n_{-1}\right]\gamma(n_{-1})\,dH(n_{-1}) \\ &= \frac{1}{2}\bar{\gamma}(X(n))\,\mathcal{G}_1[X(n),n]\,\gamma(n)\end{aligned} \tag{53}$$

It follows that $\partial H(n)/\partial C \to 0$ as $C \to 0$, as conjectured. ■

Proof of Corollary 3. The formal proof is in progress. Here we provide a heuristic sense of why the result works.

From equation (35), we can express the gross return from adjusting as

$$\Delta(n_{-1},x) = (1-\alpha)(px)^{\frac{1}{1-\alpha}}(\alpha/w)^{\frac{\alpha}{1-\alpha}} - pxn_{-1}^{\alpha} + wn_{-1} = cR(x). \tag{54}$$

A second-order Taylor expansion of $\Delta(n_{-1},x)$ around $n_{-1} = X^{-1}(x)$ yields

$$\Delta(n_{-1},x) \approx \frac{1}{2}\alpha(1-\alpha)R(x)\left[\log n_{-1} - \log X^{-1}(x)\right]^2. \tag{55}$$

Setting the latter equal to the adjustment cost, $cR(x)$, and solving implies the inverse triggers are given by

$$L^{-1}(x) \approx X^{-1}(x) \exp\left(\gamma\sqrt{C}\right), \text{ and } U^{-1}(x) \approx X^{-1}(x) \exp\left(-\gamma\sqrt{C}\right), \tag{56}$$

where $\gamma \equiv \sqrt{2/[\alpha(1-\alpha)]}$. It follows that

$$\begin{aligned} U^{-1}X(n) &\approx n\exp\left(-\gamma\sqrt{C}\right) \approx X^{-1}L(n), \text{ and} \\ L^{-1}X(n) &\approx n\exp\left(\gamma\sqrt{C}\right) \approx X^{-1}U(n). \end{aligned} \tag{57}$$

Thus, instead of arithmetic symmetry of the triggers, there is geometric symmetry. But it amounts to the same for approximate invariance of steady-state aggregate outcomes. ■

Proof of Proposition 5. The mass of firms with employment below some level n this period is equal to the mass below n in the previous period plus inflows into the mass less outflows from the mass. Thus, using equations (37) and (38) we can express the evolution of the distribution function $H(n)$ as

$$\begin{aligned} \Delta H(n)' &= G[X(n)] - H(n) - \int_{U^{-1}X(n)}^{L^{-1}X(n)} \mathcal{G}[X(n), n_{-1}]\, dH(n_{-1}) \\ &+ \int_{n}^{L^{-1}X(n)} \mathcal{G}[L(n_{-1}), n_{-1}]\, dH(n_{-1}) + \int_{U^{-1}X(n)}^{n} \mathcal{G}[U(n_{-1}), n_{-1}]\, dH(n_{-1}) \end{aligned} \tag{58}$$

Differentiating and manipulating the expression following the proof of Proposition 4 above yields the stated result:

$$\begin{aligned} \Delta h(n)' &= h^*(n)\left(1 - \mathcal{H}\left[L^{-1}X(n), X(n)\right] + \mathcal{H}\left[U^{-1}X(n), X(n)\right]\right) \\ &- \left(1 - \mathcal{G}[U(n), n] + \mathcal{G}[L(n), n]\right) h(n). \end{aligned} \tag{59}$$

Now consider a first-order approximation to $\Delta H(n)'$ around $C = 0$,

$$\Delta H(n)' \approx G[X(n)] - H(n) + \left.\frac{\partial \Delta H(n)'}{\partial C}\right|_{C=0} \cdot C. \tag{60}$$

The intercept of this expression is just the frictionless law of motion for the distribution of employment: Given any initial condition, $H(n)$, the distribution jumps to its frictionless steady state, $G[X(n)]$. We wish to evaluate the marginal effect of the adjustment cost $\partial \Delta H(n)/\partial C$ at $C = 0$.

From the proof of Proposition 4, we know that $\partial \mathcal{G}/\partial C \to 0$ as $C \to 0$. We conjecture that $\partial H(n)/\partial C|_{C=0} = 0$ and show that this implies $\partial \Delta H(n)/\partial C|_{C=0} = 0$, and thus that $H(n)$ is independent of C along the transition path. Under the conjecture, in the neighborhood of

$C = 0$ we can write

$$\frac{\partial \Delta H(n)'}{\partial C} \approx \int_n^{L^{-1}X(n)} \mathcal{G}_1\left[L(n_{-1}), n_{-1}\right] \frac{\partial L(n_{-1})}{\partial C} dH(n_{-1}) + \int_{U^{-1}X(n)}^{n} \mathcal{G}_1\left[U(n_{-1}), n_{-1}\right] \frac{\partial U(n_{-1})}{\partial C} dH(n_{-1}). \tag{61}$$

It therefore follows from equations (51) to (53) in the proof of Proposition 4 that $\partial \Delta H(n) / \partial C|_{C=0} = 0$. ■

B Labor supply

Imagine a large household with a continuum of identical members. Let L be the size of the household and N the number of employed members. Each member i derives felicity from consumption equal to $\frac{\sigma}{\sigma-1} c_i^{\frac{\sigma-1}{\sigma}}$. If member i works, he suffers disutility denoted by ζ_i. We follow Mulligan (2001) and assume that disutilities differ across members. Thus, household utility is given by

$$\frac{\sigma}{\sigma-1} \int_0^L c_i^{\frac{\sigma-1}{\sigma}} di - \int_0^N \zeta_i di. \tag{62}$$

The household's decision problem with respect to N may be recast as the choice of a threshold, $\bar{\zeta}$, such that households with disutility in excess of $\bar{\zeta}$ are not sent to work. It follows that the employment rate is given by $\frac{N}{L} \equiv n = Z(\bar{\zeta})$, where Z is the distribution function of ζ over household members. We now restate period household utility as

$$\frac{\sigma}{\sigma-1} \int_0^L c_i^{\frac{\sigma-1}{\sigma}} di - N\mathrm{E}\left[\zeta | \zeta < \bar{\zeta}\right], \tag{63}$$

where $\mathrm{E}\left[\zeta | \zeta < \bar{\zeta}\right]$ is the average disutility among workers and N is the total number of workers. Thus, the second term represents the total disutility of work within the household.[20] We use $N = LZ(\bar{\zeta})$ to write this as

$$\frac{\sigma}{\sigma-1} \int_0^L c_i^{\frac{\sigma-1}{\sigma}} di - LG(\bar{\zeta}) \int_0^{\bar{\zeta}} \zeta \frac{g(\zeta)}{G(\bar{\zeta})} d\zeta. \tag{64}$$

[20] If one interprets disutility as a fixed attribute of a worker, then workers with high disutility generally do not work but receive the same consumption as everyone else. As a result, the non-employed are always better off. This seems unpalatable, but it is not the only interpretation of the model. Instead, assume that there is a fixed distribution of disutilities, and each worker takes an i.i.d. draw from this distribution each period. That is, there are some periods where the marginal value of time is particularly high for some workers (and so they want to remain at home), and there are other periods where the disutility of work is relatively low for those same workers. In this environment, members are equally well-off on average over their (infinite) lives.

The first-order condition is

$$\bar{\zeta} = Z^{-1}(n) = \lambda w, \tag{65}$$

where w is the real wage rate and λ is the marginal value of wealth. To obtain a closed-form solution, assume the probability density of disutility is given by

$$z(\zeta) = A\zeta^{b}, \quad A > 0,\ b \text{ a real number} \tag{66}$$

in which case,

$$n = G\left(\bar{\zeta}\right) = \int_0^{\bar{\zeta}} g(\zeta)\, d\zeta = A\frac{\bar{\zeta}^{1+b}}{1+b} \tag{67}$$

and so

$$\bar{\zeta} = Z^{-1}(n) = \left[A^{-1}(1+b)\, n\right]^{\frac{1}{1+b}}. \tag{68}$$

Thus,

$$\left[A^{-1}(1+b)\, n\right]^{\frac{1}{1+b}} = \lambda w. \tag{69}$$

Since we have a free parameter, it is convenient to set $A \equiv \frac{1+b}{a^{1+b}}$, where $0 < a$. In that case, the first-order condition becomes

$$an^{\frac{1}{1+b}} = \lambda w.$$

Note that as $b \to -1$, the Frisch elasticity of labor supply goes to zero. As $b \to \infty$, the elasticity goes to infinity.[21] Also, note that if $\sigma = \infty$, then the first-order condition for consumption implies $\lambda = 1$. In that case, the labor supply schedule simplifies further to $an^{\frac{1}{1+b}} = w$, which is of the form given in the main text. Specifically, in the main text, we set $\eta \equiv \frac{1}{1+b}$ and $\psi = aL^{-\frac{1}{1+b}}$, which leaves us with $w = \varphi_0 N^{\varphi}$.

[21] Many authors now assume a Frisch elasticity of one. That corresponds to $b = 0$ and thus to a uniform distribution of disutility.

Figure 1. An *Ss* labor demand policy

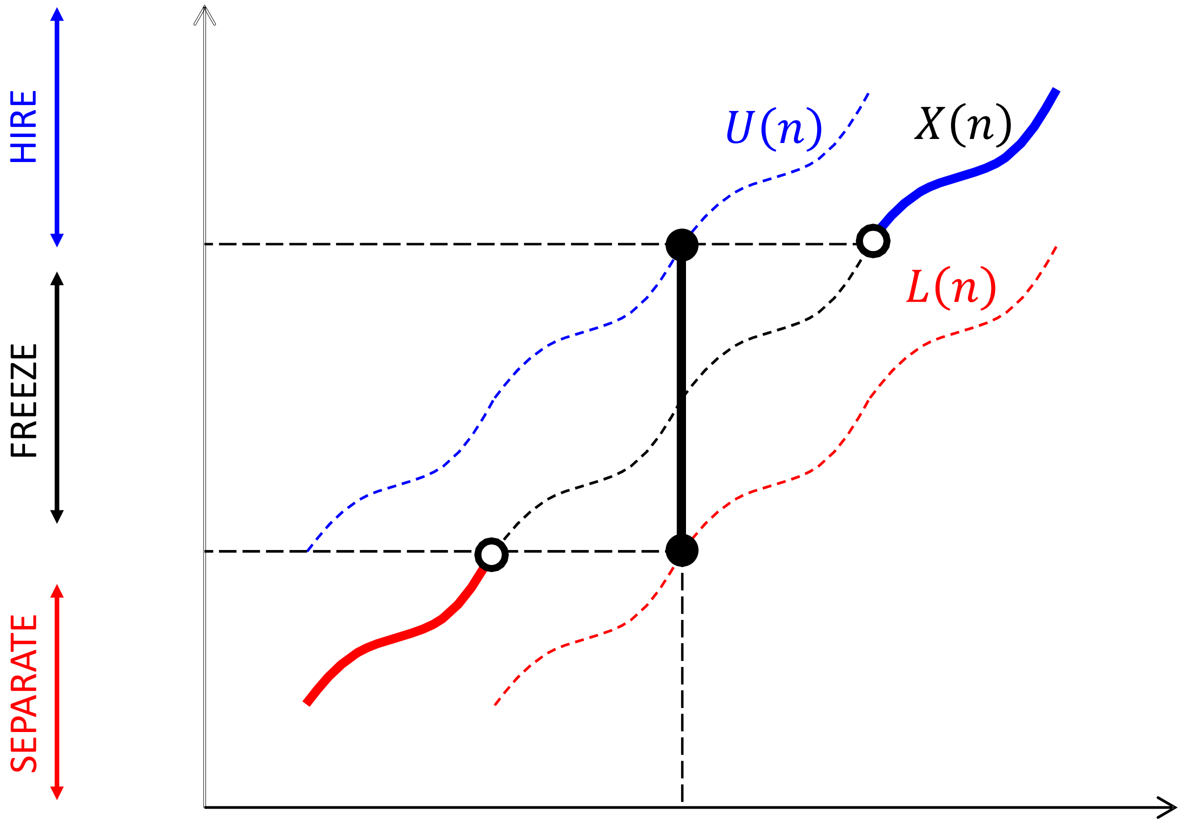

Figure 2. Intuition for near-optimality of myopia

A. Forward value as a function of current employment

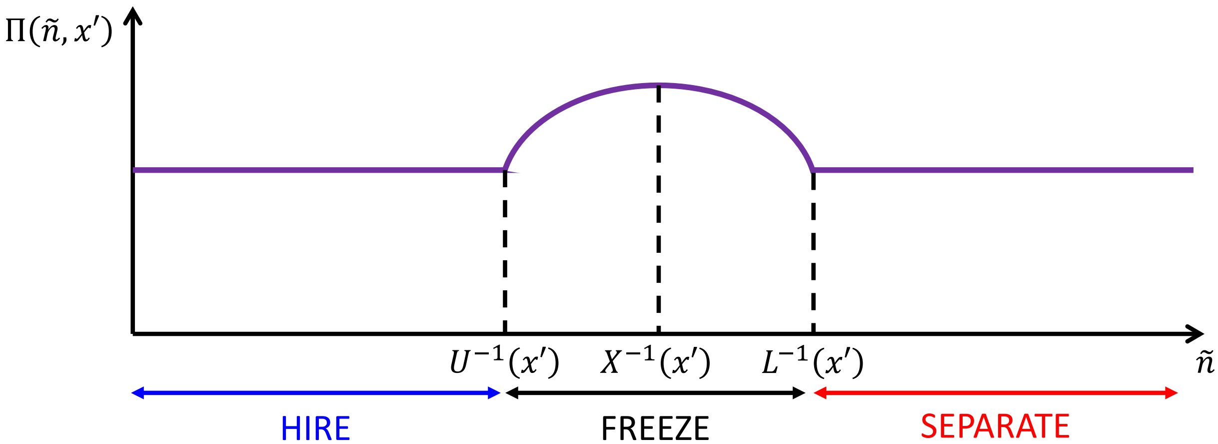

B. Marginal effect of current employment on future profits

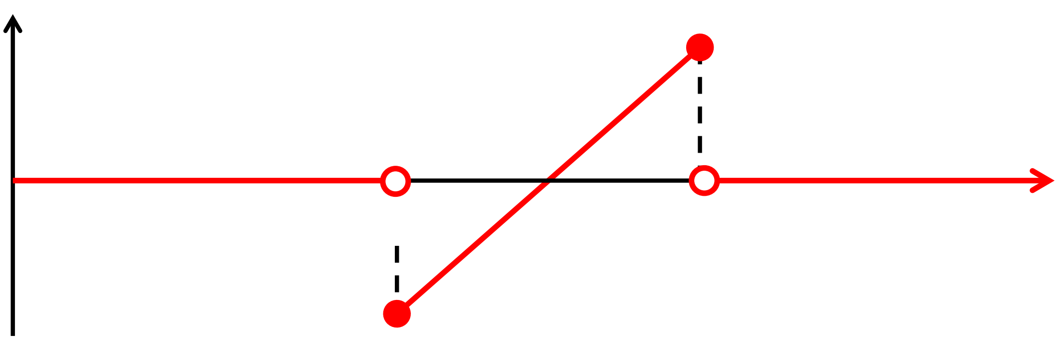

Table 1. Baseline Calibration

Parameter	Meaning	Value	Reason
α	Returns to scale	0.64	Cooper, Haltiwanger and Willis (2005)
β	Discount factor	0.99	Quarterly real interest rate = 1%
c	Adjustment Cost / Revenue	0.08	Cooper, Haltiwanger and Willis (2005); Bloom (2009)
w	Wage	0.245	Mean employment = 20
ρ_x	Persistence of x	0.4	Cooper, Haltiwanger and Willis (2005)
σ_x	Std. dev. of innovation to x	0.35	Cooper, Haltiwanger and Willis (2005, 2007)
ρ_p	Persistence of p	0.95	Autocorrelation of detrended log N
σ_p	Std. dev. of innovation to p	0.015	Std. dev. of detrended log N

Figure 3. Optimal labor demand policy: Forward-looking (**bold**) vs. myopic (shaded)

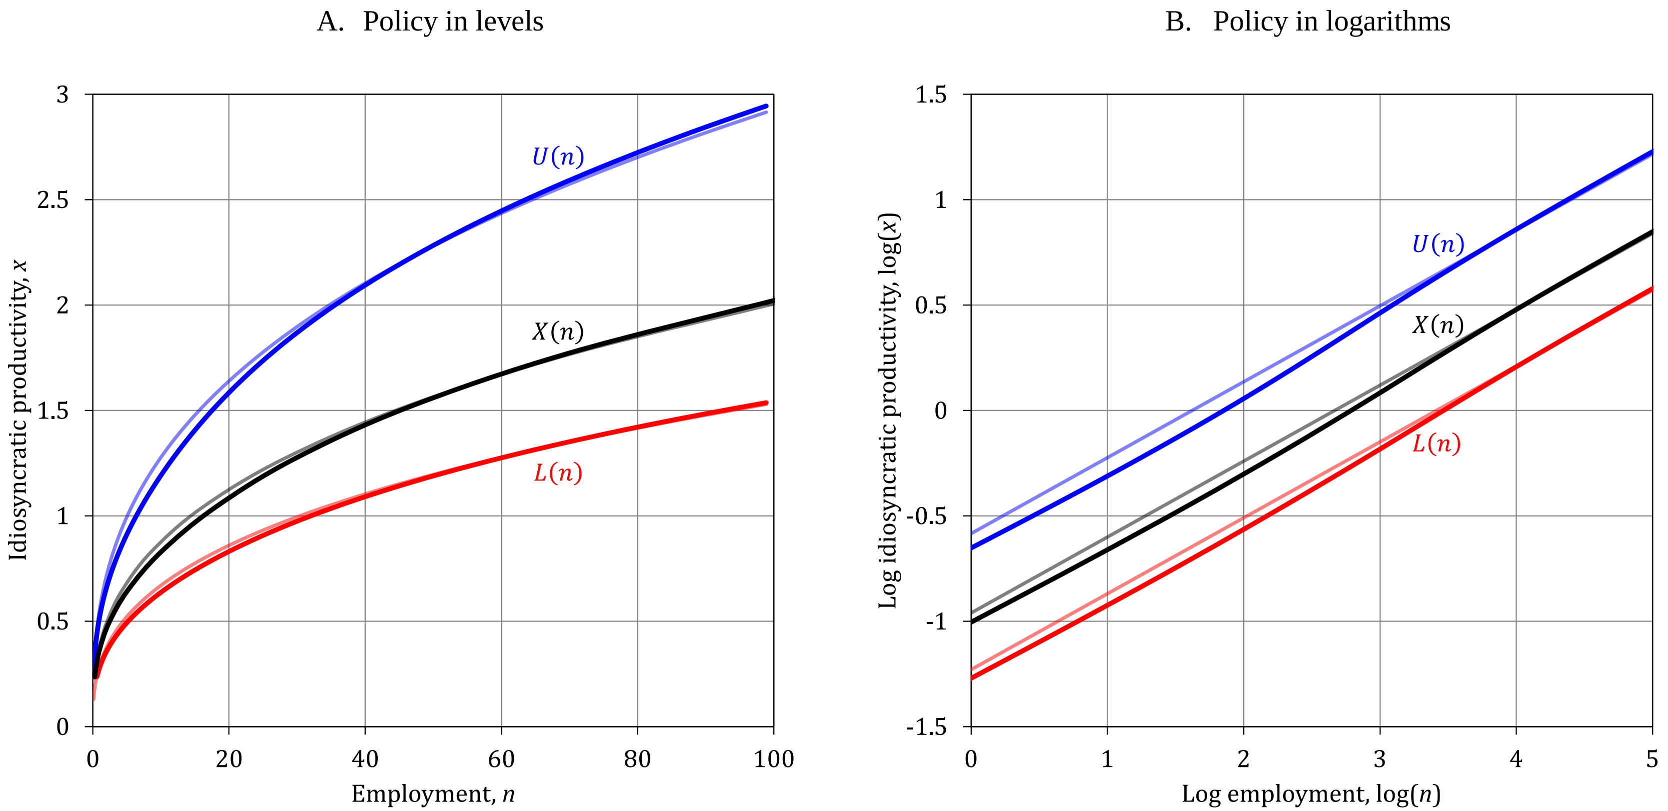

Figure 4. Employment outcomes: Forward-looking vs. myopic

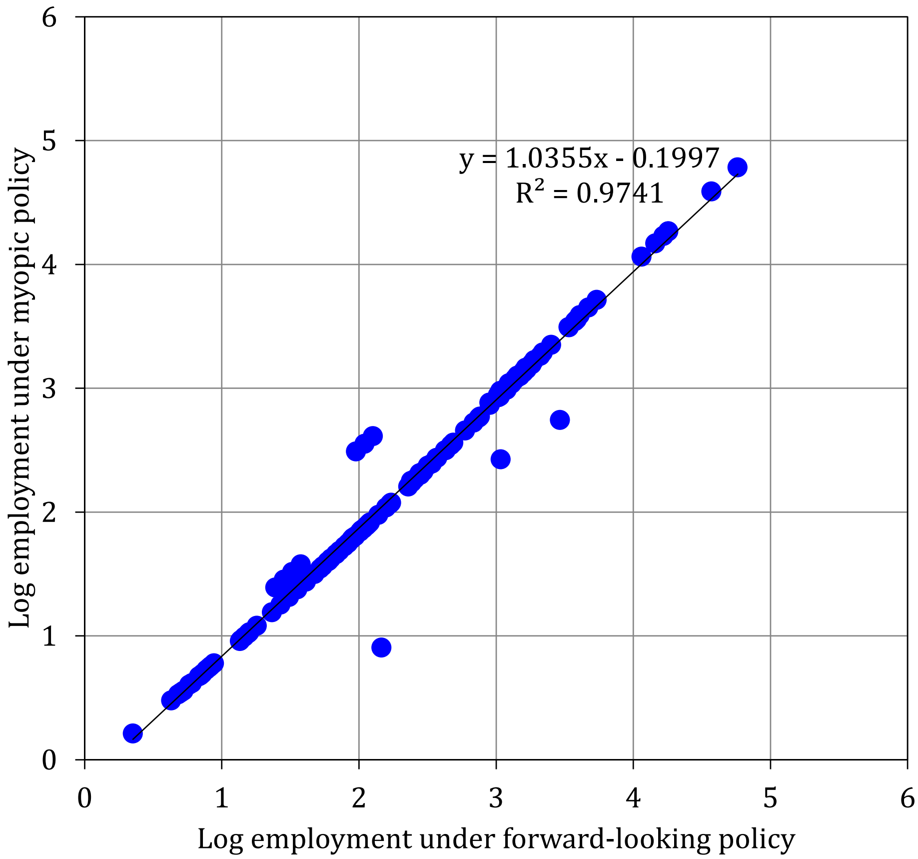

Figure 5. Intuition for the role of symmetry in aggregate neutrality

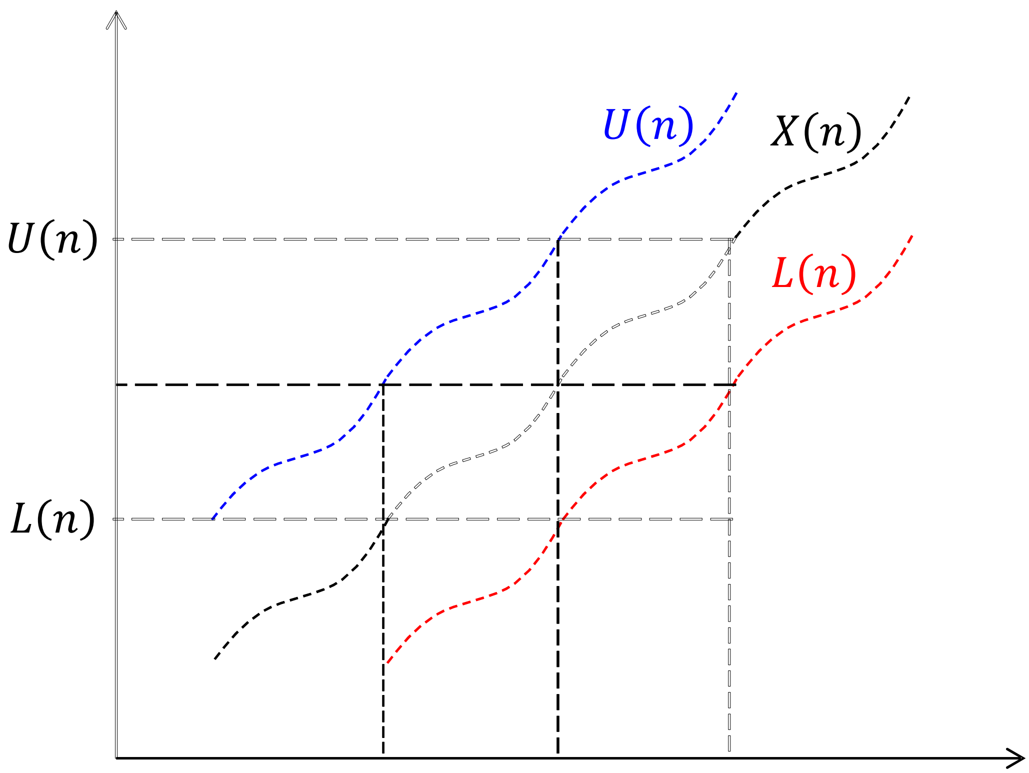

Table 2. Steady-state aggregate moments

Moment	**Forward-Looking**	**Myopic**	**Frictionless**
Aggregate (mean) employment, $N = \mathbb{E}(n)$	19.88	18.78	20.58
Cross-sectional $\sigma(\log n)$	1.007	1.064	1.058
Quarterly $\Pr(\text{adjust})$	0.462	0.452	1

Figure 6. Aggregate demand for labor

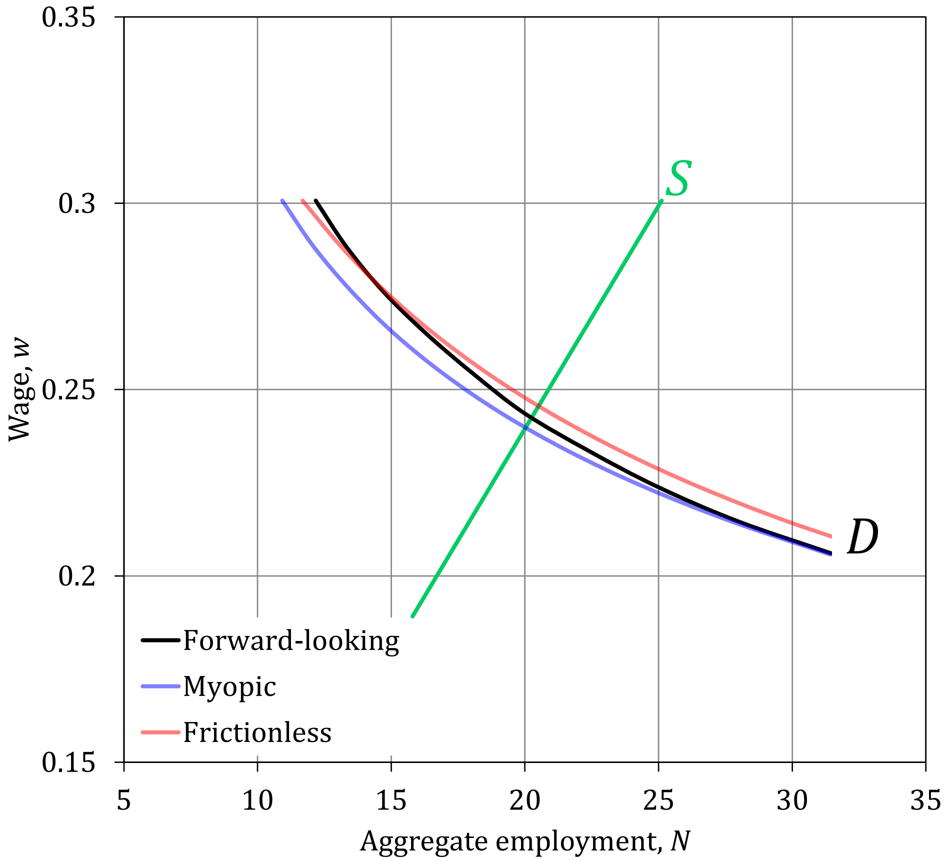

Figure 7. Dynamic response of aggregate employment to a 1% innovation to aggregate productivity