The joint dynamics of capital and employment at the plant level*

William Hawkins, Ryan Michaels, and Jiyoon Oh
University of Rochester

Preliminary and incomplete

April 11, 2012

Abstract

While there has been a wealth of research that studies the adjustment of individual factors of production, relatively little work has investigated their joint dynamics at the plant level. The present paper uses plant-level data from Chile and South Korea to investigate the joint adjustment of capital and employment. It finds a weak correlation between investment and employment growth at the plant level: although employment growth is slightly higher at an investing plant, at least one-third of establishments undertaking investment reduce their workforces, in some cases substantially. The paper argues that this fact is at odds with a model of costly multi-factor adjustment that integrates features from canonical models of dynamic capital and labor demand. The paper concludes by discussing how an extension to allow for labor-saving technological innovations can move the model closer to this moment of the data.

*We are grateful to Christopher Kurz for his assistance with the data. E-mail address for correspondence: ryan.michaels@rochester.edu.

Recent research in macroeconomics has typically approached questions regarding aggregate dynamics within models whose microeconomic environments are designed to be consistent with establishment-level observations. For instance, studies of an array of topics have incorporated adjustment costs which often make it optimal to "do nothing" in response to shocks. These Ss-type models are grounded in the empirical observation that inaction is pervasive at the micro level: in any given month or quarter, a share of establishments do not invest, change employment, replenish inventory and/or adjust product prices.[1]

Most quantitative theoretical analysis within this literature has focused on a single decision problem in order to isolate the consequences of one friction, such as the cost to invest or to change employment. All other variables under the control of a firm are assumed to be costless to adjust. There is growing interest, though, in Ss-type models which study how plants adjust along multiple margins when each choice is subject to a cost of adjusting.[2]

In reaction to this, the present paper investigates the joint adjustment of capital and employment at the establishment level. The paper exploits the fact that the integration of multiple frictions within a single model yields testable implications on the joint dynamics of the control variables at the plant level. The performance of the model along these dimensions provides valuable information as to the structure of the environment in which firms operate. These insights may then guide the re-evaluation of single-decision problems as well as the further development of models which study adjustment along multiple margins.

To organize our analysis of the establishment-level data, we consider the implications of a widely used model of capital and employment adjustment. The model assumes piece-wise linear costs of adjustment on each factor, that is, the cost of adjusting is proportional to the size of change (and the factor of proportionality may depend on the sign of the change).[3] Dixit (1997) and Eberly and van Mieghem (1997) showed that this model yields a stark prediction: whenever the relatively more costly-to-adjust factor is changed, complementarity across the factors in production implies that the less-costly-to-adjust factor is also updated.

[1] A number of papers have documented these facts on U.S. data. On investment, see Doms and Dunne (1998) and Cooper and Haltiwanger (2006); on employment, Hamermesh (1989), Davis and Haltiwanger (1992), and Cooper, Haltiwanger, and Willis (2007); on inventory, Mosser (1990) and McCarthy and Zakrajsek (2000); and on prices, Bils and Klenow (2004), Nakamura and Steinsson (2008), and Klenow and Malin (2011).

[2] See, for instance, Bloom (2009), Bloom, Floetotto, Jaimovich, Saporta-Eksten and Terry (2011), and Reiter, Sveen, and Weinke (2009).

[3] Piece-wise linear costs generate inaction because they imply a discrete change in the marginal cost of adjusting at zero adjustment. Hence, the payoff from adjusting in response to small variations in productivity does not outweigh the cost. Linear costs of control have been studied theoretically, in the context of investment, by Abel and Eberly (1996), Cooper and Haltiwanger (2006), and Khan and Thomas (2011). Ramey and Shapiro (2001) give direct evidence on the cost of reversibility. Linear costs of adjusting employment have been considered in Bentolila and Bertola (1990) and Anderson (1993).

The addition of fixed costs of adjustment makes it difficult to formally derive testable implications, but we show numerically that this does not alter the predictions of the baseline analysis.

The data analysis reveals that the model's prediction is significantly at odds with plant-level behavior. We use establishment-level data on capital and employment from the censuses of manufacturing in Chile and South Korea. Estimates of the frequency of adjustment indicate that capital is the relatively more costly-to-adjust factor; hence, the model implies investment in one period perfectly predicts (positive) employment growth in that period. Yet in the data, investment has weak explanatory power: the distribution of employment growth conditional on investment is very similar to the *un*conditional distribution of employment growth. Although the model predicts that investment and employment contraction should not co-exist, we find that between 30 percent (Chile) and 47 percent (Korea) of establishments which undertake investment also contract employment. Furthermore, among investing plants, the declines in employment at those reducing their workforce are almost as large as the expansions in employment at those increasing their workforce.

There are a few immediate tests to run to diagnose the sensitivity of this result. First, we ask whether the result is due to lag-lead patterns in factor adjustment. But we answer in the negative: employment growth is hardly more likely either in the year in which investment is done or in the next year. Second, one may argue that larger establishments are aggregates over heterogeneous production units. In that case, if one unit contracts employment substantially while another invests, the establishment as a whole is seen to reduce employment even as it invests. However, we find that the result holds for both small and large establishments. Third, as we have annual data, perhaps the result is due to time aggregation: the establishment invests in one quarter, but contracts employment significantly later in the year. We try to address this concern by simulation analysis. Specifically, we calibrate our model at a quarterly frequency and select structural parameters to replicate the annual adjustment frequencies of individual factors. We then ask whether the model replicates the relatively weak correlation between annual investment and employment adjustment and find that it does not.

This result is likely to have implications for the development of aggregate models. Business cycles are typically characterized by a strong positive co-movement of aggregate investment and employment (Barro and King, 1984). The relatively weak association between employment growth and investment at the plant level stands in contrast to this pattern in aggregate data. Hence, there is a potentially important aggregation problem to be solved to make models consistent both with plant-level outcomes and the aggregate time series.[4]

To conclude the analysis, we investigate a modification of the baseline model which helps bring it nearer to the data. The co-existence of widespread employment contraction and investment suggests the presence of labor-saving innovations. This is omitted in the baseline model. Following the related literature, our baseline model assumes Cobb-Douglas production, in which case technology is effectively Hicks-neutral. However, there appears to be an emerging consensus that the CES structure is more realistic (see Chirinko, 2008). Generalizing to a CES production structure, one can then include labor-saving technical change in a meaningful way. Factor demand outcomes then reflect a combination of neutral shocks, which move capital and labor in the same direction, and factor-biased technology shifts, which can generate investment and employment declines.[5]

The rest of the paper proceeds as follows. Section 1 introduces the baseline model and discusses its key testable implication on the joint dynamics of capital and employment adjustment. Section 2 compares the model's prediction to the data. This section discusses a number of robustness tests and, in particular, analyzes at length the implications of time aggregation. Section 3 then considers an extension of the baseline model to the case of labor-augmenting technology. Section 4 briefly surveys the related literature and situates our work in this context. Section 5 concludes.

1 The firm's problem

1.1 A baseline model

We first consider a model of capital and labor demand in which the cost of adjusting either factor is proportional to the size of the change. Formally, the costs of adjusting capital and

[4] To interpret aggregate dynamics in multi-factor models, the current literature aggregates up, in effect, from the positive co-movement between capital and labor at the plant level. For a clear analysis along these lines, see Li (2009). This line of reasoning is challenged by the plant-level data.

[5] One may interpret neutral shocks as Hicks-neutral technology shifts or as shifts in product demand. In section 3, we model the latter. Note that, if the neutral shock has an aggregate component, this may be able to generate positive co-movement of the aggregate factors over the business cycle, even as their plant-level correlation is rather weak.

employment, respectively, are assumed to be[6]

$$\begin{aligned}
\mathcal{C}_k\left(k, k_{-1}\right) &= \begin{cases} c_k^{+}\left(k-k_{-1}\right) & \text{if } k>k_{-1} \\ c_k^{-}\left(k_{-1}-k\right) & \text{if } k<k_{-1} \end{cases} \\
\mathcal{C}_n\left(n, n_{-1}\right) &= \begin{cases} c_n^{+}\left(n-n_{-1}\right) & \text{if } n>n_{-1} \\ c_n^{-}\left(n_{-1}-n\right) & \text{if } n<n_{-1} \end{cases}.
\end{aligned} \tag{1}$$

Linear costs of control have been investigated in studies of capital or labor demand, but relatively few papers have integrated both of these adjustment frictions into the firm's problem. The cost (c_n^{+}) of expanding employment is often interpreted as the price of recruiting and training, whereas the cost (c_n^{-}) of contracting employment is generally intended to represent a statutory layoff cost. Capital decisions may be costly to reverse because of trading frictions (e.g., lemons problems, illiquidity) in the secondary market for capital goods (Abel and Eberly, 1996). In that case, c_k^{+} is interpreted as the purchase price and $-c_k^{-}$ is the resale value such that $c_k^{+} > -c_k^{-} > 0$. For concreteness, we interpret the problem along these lines. But the analysis does accommodate any linear adjustment cost structure which implies costly reversibility, i.e., c_k^{+} and c_k^{-} satisfy $c_k^{+} > -c_k^{-}$. For now, we omit fixed costs of adusting from (1), but we consider the effect of these below.

The problem of a competitive firm subject to (1) is characterized by its Bellman equation,

$$\Pi\left(k_{-1}, n_{-1}, x\right) = \max_{k,n} \left\{ \begin{array}{c} x^{1-\alpha-\beta} k^{\alpha} n^{\beta} - wn - \mathcal{C}_k\left(k, k_{-1}\right) - \mathcal{C}_n\left(n, n_{-1}\right) \\ +\mathcal{D} \int \Pi\left(k, n, x'\right) dG\left(x'|x\right) \end{array} \right\}, \tag{2}$$

where x is plant-specific productivity, w is the wage rate and $\mathcal{D}$ is the discount factor. This form of the production function is particularly tractable, but the basic analysis carries through as long as (k, n, x) are complements; the technology is Hicks-neutral; and the production function displays constant returns jointly in the triple, (k, n, x). Throughout, we assume plant-level TFP, x, follows a geometric random walk,

$$x' = x\mathrm{e}^{\varepsilon'}, \quad \varepsilon' \sim N\left(-\frac{1}{2}\sigma^2, \sigma^2\right). \tag{3}$$

We omit, momentarily, both depreciation and attrition, but we will include these features in the quantitative assessment of the model. Note that we also abstract from aggregate uncertainty. This is consistent with our current focus on the cross section rather than

[6] Throughout, a prime ($'$) indicates a next-period value, and the subscript, $_{-1}$, indicates the prior period's value.

aggregate fluctuations.

Figure 1 summarizes the optimal policy and is taken from Dixit's (1997) analysis of the problem.[7] Because the problem is linearly homogeneous in $(k, k_{-1}, n, n_{-1}, x)$, the model admits a normalization.[8] We normalize with respect to x. So let us set

$$\tilde{n} \equiv n_{-1}/x, \quad \tilde{k} \equiv k_{-1}/x.$$

The figure then places $\log \tilde{n}$ along the vertical axis and $\log \tilde{k}$ along the horizontal. We summarize the policy rule with respect to employment; the capital demand rule follows by symmetry.[9] Holding k_{-1} and x constant, a higher start-of-period level of employment, n_{-1}, is tolerated within a range because of the cost of adjusting. But if the firm inherits a n_{-1} from last period that is sufficiently high, the marginal value of the worker, evaluated at n_{-1}, is so low as to make firing optimal. That is, if we let

$$\tilde{\Pi}\left(k_{-1}, n_{-1}, x\right) \equiv x^{1-\alpha-\beta} k_{-1}^{\alpha} n_{-1}^{\beta} - w n_{-1} + \mathcal{D} \int \Pi\left(k_{-1}, n_{-1}, x'\right) dG\left(x'|x\right),$$

then $\tilde{\Pi}_n\left(k_{-1}, n_{-1}, x\right) < -c_n^-$. At this point, the firm reduces employment to the point where n satisfies the first-order condition, $\tilde{\Pi}_n\left(k_{-1}, n, x\right) = -c_n^-$. Thus, the upper barrier (the northernmost horizontal line in the parallelogram) traces the values of $\log \tilde{n}$ which satisfy this FOC, making the firm just indifferent between firing one more worker and "doing nothing". Conversely, if the firm inherits an especially low value of employment, then it is optimal to hire (i.e., $\tilde{\Pi}_n\left(k, n_{-1}, x\right) > c_n^+$). Employment is then reset along the lower barrier (the southernmost horizontal line). This lower threshold thus traces the values of $\log \tilde{n}$ which satisfy the FOC for hires, making the firm indifferent between inaction and hiring one more worker.

It is important to note that, if $\log \tilde{k}$ increases, then the firm tolerates higher employment than otherwise – that is, the upper (firing) barrier is increasing in $\log \tilde{k}$. This is because of the complementarity between capital and labor. Complementarity also implies that the lower threshold is increasing in $\log \tilde{k}$; the firm is willing to hire given higher values of $\log \tilde{n}$

[7] See also Eberly and van Mieghem (1997).

[8] The assumption of a random walk is valuable analytically insofar as it preserves homogeneity. But it is also empirically plausible. Foster, Haltiwanger, and Syverson (2008) find that plant-level total factor productivity is very persistent, with an implied quarterly autocorrelation of around 0.95.

[9] The form of the optimal policy follows from the (joint) concavity and supermodularity of the value function. These properties may, in turn, be derived by a straightforward recursive argument. See Dixit (1997) and Eberly and van Mieghem (1997) for details. For the sake of brevity, our discussion is deliberately informal.

if its capital stock is larger.[10]

Figure 2, also from Dixit (1997), distills the implications of the optimal policy for the joint dynamics of capital and employment. Assume a firm has initial levels of capital and labor such that $\log \tilde{k}$ and $\log \tilde{n}$ lie in the middle of the inaction region. Now suppose productivity, x, rises, in which case $\log \tilde{k}$ and $\log \tilde{n}$ each begin to fall toward their lower barriers.[11] The figure is drawn to convey that the cost of adjusting capital is relatively high – the space between the capital adjustment barriers exceeds that between the labor adjustment barriers – so as productivity increases, the hiring barrier is the first to be reached. At this point, employment, n, is set such that $\tilde{\Pi}_n(k_{-1}, n, x) = c_n^+$.

If x continues to rise, the firm repeatedly hires in observance of its first-order condition. This implies that $\log \tilde{n}$ moves southwest along the lower barrier. As a result, when $\log \tilde{k}$ eventually reaches the investment barrier, the firm is already just indifferent between hiring and inaction. Therefore, complementarity implies that the increase in capital must tip the marginal value of labor, $\tilde{\Pi}_n$, above the marginal cost, c_n^+: *an increase in capital is always accompanied by hiring.*

A symmetric logic holds if we now suppose that x begins to decline. In this case, $\log \tilde{n}$ and $\log \tilde{k}$ reverse course and travel northeast through the parallelogram. The firing barrier is reached first, at which point the firm sets n to satisfy the first-order condition, $\tilde{\Pi}_n(k_{-1}, n, x) = -c_n^-$. When the pair $\left(\log \tilde{k}, \log \tilde{n}\right)$ later reaches the disinvestment barrier, complementarity implies that both factors are reduced. In general, then, when the relatively more costly-to-adjust factor (capital in this case) is changed, one should also witness a change, in the same direction, of the relatively less costly-to-adjust factor (employment).[12]

[10] The thresholds are flat in regions where the firm adjusts both factors. For instance, if $\tilde{n}$ and $\tilde{k}$ are sufficiently low, then the firm increases both such that $\tilde{\Pi}_k(k, n, x) = c_k^+$ and $\tilde{\Pi}_n(k, n, x) = c_n^+$. For any x, this system of first-order conditions yields a unique solution for n/x and k/x. On the figure, this unique pair is given by the southwestern corner of the parallelogram. Regardless of the exact levels of capital and employment, the firm resets to this point as long as $\left(\tilde{n}, \tilde{k}\right)$ initially lies to the southwest. Hence, in this region, the hiring threshold is independent of the initial level of capital and the investing barrier is independent of the initial level of employment.

[11] If x rises by one log point, for instance, then $\tilde{n}$ and $\tilde{k}$ each fall by one log point. Hence, the pair $\left(\log \tilde{n}, \log \tilde{k}\right)$ travels along the 45^0 degree line. This simple characterization is made possible when both factors are expressed in logs, which explains why we do so in Figures 1 and 2.

[12] We have argued that there is a certain ordering of adjustment by simply tracing the pair $\left(\tilde{n}, \tilde{k}\right)$ through Figure 2. Eberly and van Mieghem (1997) formalize this claim (see their Proposition 3). The key step is to establish that the slopes of the capital adjustment barriers in Figure 2 exceed one. They show that this follows from homogeneity of the value function. As a result, if the pair $\left(\tilde{n}, \tilde{k}\right)$ begins at the southwestern-most point and traces a path northeast along the 45^0 line, it will always lie south of the investment barrier. Accordingly, the next barrier it will reach is the firing threshold. An analogous argument applies if $\left(\tilde{n}, \tilde{k}\right)$

1.2 The model with fixed adjustment costs

One may generalize (1) to include fixed costs of investing $\left(C_k^+\right)$, disinvesting $\left(C_k^-\right)$, hiring (C_n^+), and firing (C_n^-):

$$\begin{aligned} \mathcal{C}_k\left(k, k_{-1}\right) &= \begin{cases} C_k^+ k_{-1} + c_k^+\left(k - k_{-1}\right) & \text{if } k > k_{-1} \\ C_k^- k_{-1} + c_k^-\left(k_{-1} - k\right) & \text{if } k < k_{-1} \end{cases} \\ \mathcal{C}_n\left(n, n_{-1}\right) &= \begin{cases} C_n^+ n_{-1} + c_n^+\left(n - n_{-1}\right) & \text{if } n > n_{-1} \\ C_n^- n_{-1} + c_n^-\left(n_{-1} - n\right) & \text{if } n < n_{-1} \end{cases}. \end{aligned} \tag{4}$$

In quantitative applications, it is common to scale the fixed cost, so that large, productive firms do not "outgrow" it. This is what we do here. Specifically, we will assume the fixed cost of adjusting capital (employment) is proportional to the start-of-period stock (workforce).

An analytical characterization of the policy is difficult in this case, but the factor demands can be solved numerically. Figure 3 plots the solution based on a calibrated version of the model (2)-(4).[13] (Ignore the arrows on the figure for the moment.) As is well known, the policy rule consists of two sets of functions. The inner parallelogram (solid lines) traces the reset policy: if it adjusts, the firm resets capital and employment to a point along these schedules. These schedules thus give the pairs of $\left(\tilde{n}, \tilde{k}\right)$ such that, conditional on adjusting, the marginal value of a factor equals the marginal cost of adjusting that factor. The outer hexagon (marked by "$\times$") is formed by the triggers that determine the choice to adjust. For instance, on the far left is the investment trigger, which gives, for any initial (normalized) level of employment $\tilde{n}$, the largest level of initial (normalized) capital such that the firm prefers to invest for all capital less than the trigger and prefers not to invest for all levels greater than the trigger. Analogous definitions hold for the other triggers. Note that the presence of a fixed cost implies economies of scale in adjusting, so it is optimal to wait longer than otherwise between instances of adjustment and then to make large, or lumpy, changes. Hence, the investment triggers lies to the west of the reset schedule: the firm permits $\tilde{k}$ to drift beyond the point at which it would exercise regulation in the absence of fixed costs.[14]

---

sits at the northeastern-most point of the parallelogram. Thus, capital is only adjusted when $\left(\tilde{n}, \tilde{k}\right)$ reaches the southwestern or northeastern tips of the parallelogram in Figure 2.

[13] The details of the calibration are given in the next section, where we also discuss the model's aggregate steady-state properties. For now, we just note that the structural parameters are selected to roughly replicate the frequencies of adjustment of each individual factor.

[14] The reader may notice that the trigger on investment does not increase monotonically in $\tilde{n}$. At first, this seems counterintuitive: by complementarity, a lower level of employment reduces the incentive to invest, so one may expect the investment trigger to decline everywhere in $\tilde{n}$. This does not appear to happen in the southwest corner of the diagram. The reason is that the firm is nearing its hiring trigger in this region. As the probability of a large, discrete rise in employment is high, complementarity implies that the firm

The arrows on Figure 3 illustrates the pattern of adjustment under fixed costs. Specifically, they trace the path of $\left(\tilde{n},\tilde{k}\right)$ in response to a series of positive productivity innovations. Over the first several periods, a series of upward adjustments are made to employment, with no changes in capital. Then, at the most southwestern tip (where the arrow stops), the plant increases both capital and employment to the point identified by the circle in figure. If x then began to fall, the pair $\left(\tilde{n},\tilde{k}\right)$ would trace a path northeast until it reached the firing trigger. Hence, the only difference in the dynamics induced by fixed costs is that factor demand changes, conditional on adjusting, are lumpy. But as long as the total costs of adjusting employment remain relatively small, the prediction regarding the joint dynamics of the two factors is unaltered: adjustments to capital always coincide with employment changes.

The dynamics are driven, again, by complementarity. To recall, in the absence of fixed costs, complementarity implies that capital adjustment tips the *marginal* value of labor above the marginal cost. In the presence of fixed costs, it is not enough that the marginal value of a new hire exceed c_n^+. Instead, because of the fixed, discrete adjustment cost, it must be that the *total* value of adjusting to n $\left(\text{i.e., } \tilde{\Pi}\left(k_{-1},n,x\right)\right)$ exceeds the value of inaction. By complementarity, large, or lumpy, changes in capital raise the former substantially and thus induce employment adjustment.[15]

2 Model evaluation

2.1 An initial assessment

To assess the implications of the model, we use two sources of plant-level data. The first is the Chilean Manufacturing Census, for which we have annual data from 1979-96. The Census surveys all manufacturing establishments with at least 10 workers. The data has been widely used in the trade and industrial organization literature.[16] Our second source is the Korean Annual Manufacturing Survey, for which we have data from 1990-2006. This, too, is a census of manufacturing plants with at least 10 workers.[17] Both surveys include observations on investment and the size of the plant's workforce. In what follows, we focus

becomes more inclined to invest.

[15]We should stress that our analysis pertains to the employment decision observed at the instant of investment. In annual data, however, we do not observe the instant of investment – firms make decisions more often than once per year. Thus, as mentioned in the introduction, time aggregation over multiple decisions through the year can yield (year-over-year net) employment declines coincident with investment. In the next section, we discuss time aggregation in detail.

[16]See Liu (1993), Levinsohn (1999), and Petrin and Sivadasan (forthcoming).

[17]Beginning in 2006, plants with between 5 and 10 workers were included in the survey. For comparability across time (and with respect to Chilean data), we therefore stop the sample at 2006.

on equipment investment specifically; preliminary analysis reveals that the conclusions holds for structures.

One challenge with the Chilean data should be mentioned at the outset. Employment in the survey is typically measured as an annual average. As a result, the measured change in employment does not necessarily accurately reflect the change in employment from the end of one (calendar) year to the end of the next. In principle, this complicates the comparison of employment growth to investment, since the latter is measured as the cumulation of capital goods purchases (net sales) over a calendar year. Fortunately, though, there were several consecutive years in the 1990s in which the Census also included questions about employment as of the end of the year. When we restrict the sample to use point-in-time data from these years, our results are unaffected. Moreover, our Korean data measures end-of-year employment consistently, and reveals the same pattern in terms of the joint adjustment of capital and employment. Thus, we are doubtful that this particular measurement issue has significantly influenced our findings.

We find that capital appears to be the more costly to adjust factor in both Chilean and Korean data. Indeed, the adjustment probabilities are remarkably similar across the two datasets. In both, almost half of all plant-year observations involve zero equipment investment. In contrast, around 15% of plant-year observations involve zero (net) employment adjustment.

In light of this fact, we ask whether investment (or disinvestment) is always accompanied by employment adjustment. Figure 4 shows that it is not. The figure plots the unconditional distribution of the change in log employment and the distribution conditional on plant-level investment in excess of 10 percent (that is, investment relative to start-of-period capital exceeds 0.1).[18] We use the 10 percent threshold for the investment rate because we are more concerned that meaurement error in investment may mean that some of the smaller reported investments are in fact zeros. If the model of section 1 is right, this would bias the correlation of employment growth conditional on investment toward zero, and thus lead us to overstate our result.

The employment growth distribution, conditional on positive investment, is slightly shifted to the right – plants do typically raise employment more if they also invest – but the similarity across these distributions is inconsistent with the model. The model predicts

[18]Whereas the Korean survey provides annual data on a plant's capital, the Chilean Census gives an estimate of the capital stock only in 1980 and 1981 and later, in the years from 1992-96. To fill in the missing years, we use the perpetual inventory method, as in Petrin and Sivadasan (forthcoming, Appendix D). Equipment is depreciated at a 10 percent annual rate and deflated with the industry capital goods price deflator.

that investment should perfectly predict an expansion in employment; employment growth should lie everywhere to the right of zero. Instead, we find that, in Chile, about 30 percent of the establishment-year observations which show investment also report a contraction in employment. In Korea, the share is higher, at nearly 47 percent. Moreover, the average contraction among this set of plants is fairly large: plants which contract, conditional on investment, reduce their workforce by 16.5 percent in Chile and 22 percent in Korea. Indeed, the magnitude of the average decline is not that much different than the magnitude of the average increase (conditional on investment), which is a little over 20.5 percent in Chile and 24.4 percent in Korea. (These numbers are reported in Table 2.)

Another way to illustrate the weak correlation between investment and employment growth is to simply regress the latter on the former. When we do this, we find that a one percentage-point increase in the investment rate (that is, investment relative to start-of-period capital) is associated with an increase in employment growth of roughly one-tenth of one percentage point. As we will see, our quantitative analysis of the model of section 1.1 reveals that the model implies a coefficient that is almost an order of magnitude higher.

We have undertaken a number of robustness tests. For instance, perhaps there is a time required to install machines so workers are hired with a lag, after the equipment has been put in place. This suggests that one should consider the relationship between investment in one year and employment in the next. But the kernel densities of log employment changes conditional on lagged investment (in excess of 10 percent) look virtually the same as in Figure 4. In fact, employment contraction, conditional on past investment, is slightly more likely.

Second, one may suspect that Figure 4 is partly due to aggregation over heterogeneous production units within large establishments. Perhaps one division of the establishment undertakes investment and hires. Another division contracts employment substantially, though does not disinvest. The net establishment-wide employment change may well be negative, even though establishment-wide investment is positive. If this is right, we should see different joint adjustment dynamics across size classes. Yet this is not true. It remains the case that, if we restrict the Chilean sample to establishments with fewer than 25 workers, 34.6 percent of the establishment-year observations for which we see investment also report a decline in employment. The share in Korean data is slightly higher, at over 40 percent. Moreover, in both datasets, the average contraction is roughly the same as the average expansion (conditional on investment).[19]

[19] That we are able to repeat the analysis on relatively smaller plants is an important advantage of these data. This sort of robustness test would not be possible if we performed this analysis on Compustat data,

Third, it is possible that Figure 4 reflects time aggregation. Establishments may invest and hire in one month but later reduce employment more significantly. As a result, the data show positive investment over the year but a net employment decline. This concern is more difficult to address. In the next subsection, we detail our approach.

2.2 Robustness to time aggregation

Given the available data, it is impossible to investigate time aggregation directly, of course. We explore this within the context of a calibrated version of the model presented in section 1.1. The model is calibrated at a quarterly frequency. The intention is to simulate quarterly plant-level decisions and determine whether time-aggregating these observations to an annual frequency yields the joint adjustment dynamics observed in the annual data. Admittedly, the calibration is somewhat illustrative, but the simulations are nonetheless instructive.

There are five structural parameters in particular whose values must be pinned down: the adjustment costs, $\left(c_k^+, c_k^-, c_n^+, c_n^-\right)$, and the standard deviation, σ, of the innovation to productivity. The price of the investment good, c_k^+, is normalized to one. The resale price, $-c_k^-$, is then calibrated to target the frequency of investment. The intuition behind this strategy is straightforward: a large wedge between the purchase and resale prices raises the cost of reversing investment decisions, and so induces a greater degree of inaction. Next, for simplicity, we impose symmetry on the costs of adjusting the workforce, $c_n^+ = c_n^-$. We then choose this cost to target the frequency of net employment adjustment. Lastly, the variance of the innovation is set to be sufficiently high in order to target the average expansion in (log) employment among those establishments which increase employment and undertake investment. Given this, we will ask below whether the model is also capable of generating the average *contraction* in (log) employment among those establishments which decrease employment and do investment.[20]

The other parameters are selected partly based on external information and partly to be consistent with choices in the related literature. The full list of parameters is given in Table 1. The discount factor, $\mathcal{D}$, appears somewhat low, but it is set to be consistent with

which includes only (relatively large) U.S. publicly traded corporations.

[20] The choice of c_n^- implies a much lower cost of dismissal than suggested by a literal reading of Chilean statue. During the 1980s, the estimated present discounted cost of mandated severance rose to three months of wages on average (Petrin and Sivadasan, forthcoming). Yet, to replicate the inaction in the employment growth distribution, we must set c_n^- to be only around one third of the monthly wage. In principle, one could set c_n^+ to zero, in which case inaction arises only because of the cost of dismissal; this would force a higher calibration of c_n^-. However, this approach would also imply a great deal of skewness in the employment growth distribution, with the left tail truncated. This would be inconsistent with the rough symmetry in the empirical distribution.

the average real interest rates in Chile and South Korea over the time periods for which we have data. The elasticity of output with respect to labor, β, is set to be consistent with a labor share of around 55 percent in Chile (Kumhof and Laxton, 2009) and 60 percent in South Korea (Oh, 2011). To guarantee a well-defined notion of plant size, the coefficient, α, attached to capital must then be set below $1-\beta$. Without much guidance on this, we fix $\alpha = 0.25$.[21]

Once the model is calibrated, it is solved via value function iteration. The homogeneity of the value function with respect to x is helpful at this point, as it allows us to re-cast the model in terms of the normalized variables, $\tilde{n}$ and $\tilde{k}$. This eliminates a state variable. Once the policy functions are obtained, we simulate 20,000 plants for 250 quarters. Results are reported based on the final 20 years of data.

Table 2 summarizes the simulation results for the baseline model. The table reports a few sets of statistics, each computed off the simulated panel. The top panel reports moments for individual variables and is designed to gauge how the model matches the marginal distributions of each factor change. We report, for instance, the inaction rates with respect to employment and investment, shown in the top two rows. (More precisely, these are the shares of establishment-year observations for which there is no net change in employment and no investment.) These were moments targeted in the calibration, so the model (nearly) reproduces the empirical estimates.

In the top panel, we also report a measure of the dispersion in employment growth, taken here to be the unconditional standard deviation of the year-over-year log change. This moment was not targeted, and the model does noticeably understate dispersion. An increase in σ would ameliorate this, but the presence of larger shocks would also imply larger average employment changes among expanding plants. Given the present calibration, the model already overstates this latter moment (see the penultimate row of the table). Our choice of σ tries to balance these competing concerns, to some extent.[22]

Lastly, in the top panel, we see that the model does not replicate the skewness in the empirical distribution of annual investment. In the data, between 3.5 (Chile) and 5.75 (Korea)

[21] The elasticities of output with respect to labor are set to be indicative of the production technology within the macroeconomies of these two countries. However, the labor shares within our datasets, which include only manufacturing plants, are notably lower. It is unclear why this calibration would yield qualitatively different results with regard to the joint adjustment of the two factors, but a robustness analysis is nonetheless underway.

[22] An increase in σ would tend to make time aggregation more salient, all else equal: larger shocks raise the probability that a plant invests but later in the year fires. However, to replicate the inaction rates on each factor, one could not raise σ and leave all else constant. Instead, the costs of adjusting would also have be increased, which would in turn restore a relatively large region of inaction.

percent of establishment-year observations involve disinvestment, whereas the model-implied distribution is more symmetric about zero. As we show below, though, the introduction of depreciation vastly improves the model's performance along this dimension.

The next panel summarizes the joint adjustment of capital and employment. The first three rows represent the cleanest check on the analysis in section 1. These report moments at a *quarterly* frequency, so time aggregation is not a concern. The rows report (i) the share of establishment-quarter observations for which investment is observed which also involve a net reduction in employment; (ii) the average decline in (log) employment among those establishment-quarter observations which involve contraction and positive investment; and (iii) the average increase in (log) employment among those establishment-quarter observations which involve expansion and positive investment. As the analysis of section 1 predicts, positive investment and employment contraction never jointly occur.

The next few rows speak to the issue of time aggregation. The model implies that time aggregation does not account for the empirical results. The share of establishment-year observations which involve both positive investment and employment contraction is still virtually zero. In addition, conditional on contracting given positive investment, the average decline in employment is less than 1/8 the size in the data. These results are shown graphically in Figure 5, which plots the model-implied distribution of the log changes in employment given positive investment (in excess of 10 percent). For reference, we show the distribution from the Chilean data next to it.

2.2.1 The impact of time aggregation in the presence of depreciation

However, time aggregation may take on greater significance under slightly different circumstances. Suppose, for instance, that we introduce geometric depreciation into the model of section 1.1.[23] To get some insight into the dynamics in such an environment, we return to Figure 2. Depreciation affects the angle at which the pair $\left(\log \tilde{k}, \log \tilde{n}\right)$ travels through the space: as x increases, $\log \tilde{k}$ falls more than otherwise, so the angle becomes more shallow. This means that $\left(\log \tilde{k}, \log \tilde{n}\right)$ approaches a given investment barrier at a faster rate. Conversely, as x falls, the pair $\left(\log \tilde{k}, \log \tilde{n}\right)$ climbs at a sharper angle toward the firing barrier, because depreciation limits the rise in $\tilde{k}$. This "bias" toward firing is intuitive: since depre-

[23] Formally, it is straightforward to incorporate constant geometric depreciation and attrition. Assume capital depreciates at rate, δ_k, and a firm's workforce attrites rate, δ_n. Since k_{-1} and n_{-1} are the end-of-last-period values, the firm carries into the present period a stock of capital equal to $\hat{k}_{-1} \equiv (1-\delta_k)\, k_{-1}$ and a workforce of size $\hat{n}_{-1} \equiv (1-\delta_n)\, n_{-1}$. In (1) and (2), then, we simply replace k_{-1} with $\hat{k}_{-1}$ and n_{-1} with $\hat{n}_{-1}$, and the analysis proceeds as before.

ciation erases some of the undesired capital but the firm has the same number of workers, there is more "excess" employment than in the no-depreciation baseline.

The effect of this "bias" toward investing and firing on the joint dynamics is slightly subtle. On the one hand, if depreciation is not too large and the costs of adjusting capital remain sufficiently high, the dynamics, at a *quarterly* frequency, should be the same as in the model of section 1.1. Under these conditions, the pair $\left(\log \tilde{k}, \log \tilde{n}\right)$ will still reach the barriers on employment first, and investment will continue to perfectly predict employment growth at that instant.[24] On the other hand, the bias toward investment and firing does suggest that the effect of time aggregation on the *annual* moments may be more pronounced in the presence of depreciation. After investment, the pair $\left(\log \tilde{k}, \log \tilde{n}\right)$ is more likely than in the no-depreciation baseline to reach the firing barrier over the subsequent quarters.

Table 3 (first column) reports the same moments as Table 2 for the model with depreciation. We set the depreciation rate to be 10 percent per year, which is the rate used by researchers to actually calculate capital in the Chilean data via the perpetual inventory method (see footnote 18). A few results are noteworthy. First, the model now largely replicates the empirical distribution of investment (precisely, the level of investment relative to the prior year's capital). The frequency of disinvestment, for instance, is much nearer to the data. This reflects the fact that, since disinvestment is costly, a plant allows depreciation to erode units of unwanted capital. Hence, it disinvests much less often. If we also compute the standard deviation of investment (not shown in the table), it is 0.2 in the data versus 0.215 in the simulated data

Second, as for the joint dynamics, it remains the case that, at a quarterly frequency, positive investment hardly ever coincides with employment contraction. Furthermore, though the share of observations in the simulated data which involve both annual net employment declines and investment is higher than in the baseline model, it is no greater than 1/6 of what is observed in the data. The average decline among contracting establishments, conditional on investment exceeding 10 percent, is also counterfactually low.

At the same time, we should note that the simulation results are sensitive to the choice of the threshold for positive investment. If we condition instead on *any* positive investment (rather than at least 10 percent), over 25 percent of observations in the simulated data which involve any positive investment also show a (year-over-year net) employment decline. This is

[24] Moreover, one should note that the capital barriers will not remain fixed after the introduction of depreciation. A forward-looking firm, anticipating depreciation, will both invest sooner than otherwise and defer disinvestment to a greater degree, instead allowing depreciation to erode undesired capital. Hence, the capital barriers shift to the left, partially blunting the effect of depreciation (whatever its size) on the time it takes $\tilde{k}$ to reach the investment barrier.

near the empirical analogue for this moment in Chilean data, which is around 36 percent.[25] That said, it remains the case that the average decline among contracting establishments, conditional on any positive investment, is counterfactually low: it is -5.8 percent in the simulated data, but 18 percent in plant-level data.

Hence, in the model, contraction does occur (on net over the year), but it is limited. This reason for this is intuitive. As suggested above, the pair $\left(\tilde{k}, \tilde{n}\right)$ will likely approach the firing barrier more quickly in the presence of depreciation. But firing only undoes the hiring that accompanied the recent investment episode if the number of hires was small. That in turn means that the increase in investment must have been relatively small, too. So episodes in the model of firing and positive investments are ones where the investment was quite limited, e.g., less than 10 percent.

This is a revealing property of the model because it contrasts so clearly with the plant-level data. The share of observations which involve investment and employment contraction in the actual data is much more robust to the choice of the investment threshold. As we raise the threshold from zero to 10 percent, for instance, this share falls modestly from 36 percent to 30 percent (in Chilean data). We can also condition on investment greater than 20 percent rather than 10 percent. In this case, in model-generated data, episodes of investment and employment contraction virtually vanish: the share of establishment-year observations that involve both falls from 5.6 percent (shown in Table 2) to 0.25 percent. In the actual data, in contrast, the share falls only slightly to 29 percent. This suggests that, even in the presence of constant depreciation, time aggregation likely does not account for our findings.

Before we leave the topic of depreciation, there is one additional test we wish to run. The depreciation process at the plant level is arguably badly described by the assumption of constant, geometric decay. This restriction, moreover, might affect our results. To see this, consider an alternative. Assume there are occasional, and at least partially unforecastable, instances of large, or lumpy, depreciation. This specification of the depreciation process reflects the fact that breakdown may in fact occur suddenly at the plant level. There may be smooth, geometric decay in the efficiency of capital from the time of installation, but the inter-operability of machine parts implies that the entire machine can break down abruptly once decay reaches a threshold. The machine *as a whole* would then have to be replaced.[26]

Large, or lumpy, depreciation gives the plant a motive to invest even if plant-level pro-

[25] Because of recent difficulties in remotely accessing the data in Korea, we have not been able to compute this statistic using Korean data.

[26] This is the sense in which machines are "lumpy" (see Cooper, Haltiwanger, and Power, 1999). The idea that service lives of capital are stochastic is an old idea in the literature on investment (see the discussion in Feldstein and Rothschild, 1974). We marry these two ideas to develop a specification for depreciation.

ductivity is relatively low (and thus labor demand is relatively low). Figure 6 helps visualize this. A large depreciation of capital shifts the pair, $\left(\tilde{k}, \tilde{n}\right)$, to the left and induces the firm to invest and reduce employment.[27]

To represent this notion of lumpy depreciation in a tractable way, we assume a modified "one-hoss-shay" depreciation process in which capital may not decay for several periods and but can then decline significantly. Specifically, we assume that the quarterly depreciation rate, δ_k, is drawn from a two-point distribution: with 80 percent probability per quarter, $\delta_k = 0$, and with 20 percent probability, $\delta_k = 0.125$.[28] This specification preserves an annual average depreciation rate of 10 percent, as in the case analyzed above. But it allows for instances in which 12.5 percent of a plant's capital fails.

This is, of course, just one of many possible calibrations of this depreciation process. However, we emphasize that the average annual rate of depreciation does discipline the size of the decline in δ_k. If one wishes to have δ_k fall instead by 20 percent, these instances of depreciation must happen much less often to preserve an average annual rate of depreciation of 2.5 percent. This suggests to us that the results are unlikely to be too sensitive to the exact calibration of the "lumpy" depreciation.

Table 3 (second column) reports the results. Consistent with our intuition, lumpy depreciation does induce a greater incidence of employment contraction and positive investment, even at a quarterly frequency. But the model-generated outcomes still remain appreciably far from their empirical counterparts. The probability that a plant invests and reduces employment is still less than half of its empirical analogue, and the average contraction, conditional on positive investment, is no more than 40 percent of its value in the data.

2.2.2 Further sensitivity analysis to time aggregation

The introduction of worker attrition is rather innocuous compared to depreciation. All else equal, attrition biases employment adjustment toward expansion. Furthermore, attrition implies that, again all else equal, the pair, $\left(\tilde{k}, \tilde{n}\right)$, approaches a given disinvestment barrier at a faster rate. This follows from the fact that $\left(\tilde{k}, \tilde{n}\right)$ will climb at a more shallow angle (relative to the no-attrition baseline), because attrition limits the increase in $\tilde{n}$. So, at first glance,

[27] One may also interpret depreciation as obsolescence. At the plant level, there is no reason to believe that this occurs smoothly. Evidence on the thinness of resale markets for machinery (see Ramey and Shapiro, 2001) suggests that a firm's capital is highly customized. This means that the law of large numbers need not hold: innovations in the production of a plant's specific class of equipment can occur suddenly and thus abruptly make its current capital obsolete.

[28] The depreciation rate, δ_k, is thus i.i.d. over time. The absence of persistence simplifies the calculation of the solution, as it removes δ_k as a state variable.

we might expect more instances, at least an annual frequency, of hiring and disinvestment. This does not particularly help us engage the data, however, as disinvestment is exceedingly rare empirically. The co-movement of hiring and positive investment is, moreover, virtually unaltered relative to the baseline model of section 1.1. The results in Table 4 (first column) illustrate this.[29]

Lastly, we ask whether the time-aggregated results are at all different once fixed costs of adjusting capital and employment are included. Our discussion in section 1.2 suggested that fixed costs are unlikely to alter the dynamics, provided the total cost of adjusting capital remains relatively high. Table 4 (second column) confirms this. The coincidence of positive investment and employment contraction is nearly non-existent.[30]

3 Extending the baseline model: Introducing labor-saving innovations

In light of the analysis of the prior section, our working hypothesis is that the canonical model of section 1 is inconsistent with the joint dynamics of capital and employment. In this section, we discuss an alternative specification of the firm's problem that appears more promising.[31]

Labor-augmenting technical change plays a prominent role within the recent growth literature (see, for instance, Acemoglu, 2001), but it has been absent from recent macroeconomic models of dynamic factor demand.[32] Yet it is a natural place to begin our investigation.

[29] We note that the probability that employment does not change year over year is virtually zero once exogenous worker attrition is included. This is a problem with this specification of attrition that is common to models of dynamic labor demand.

[30] For this exercise, we assume symmetry, so $C_k^- = C_k^+ \equiv C_k$ and $C_n^- = C_n^+ \equiv C_n$. This reduces the number of free parameters, which simplifies the calibration. Moreover, studies of dynamic factor demand typically impose this assumption, so there is no evidence to guide the calibration of asymmetric costs. The parameters, C_k and C_n, are then selected as follows. We first make the linear costs close to, but strictly greater than, zero. This serves to emphasize the role of the fixed costs, while preserving a degree of continuity with the model of section 1.1. Then, we raise C_k and C_n until a combination of adjustment costs is found which induces the empirical frequencies of capital and employment adjustment.

[31] The modifications are on the (factor) demand side; we continue to abstract from general-equilibrium considerations. This approach is motivated by an intuition that price movements in general equilibrium are unlikely to generate qualitatively different joint dynamics at the plant level. To see this, return to the frictionless model of capital and labor demand. If the revenue function displays decreasing returns jointly to the two factors, then an increase in the price of either factor (perhaps because of a sudden withdrawal of supply) reduces demand for both factors – factor demands react in the *same* direction. As a result, our intuition is that, in the frictional model, a sudden rise in, say, the wage facing a plant will likely trigger firing (and thus a higher capital-labor ratio), but it is unlikely to simultaneously trigger positive investment.

[32] Within the literature on fluctuations, labor-saving technology has made a belated appearance. Francis

Under certain conditions, labor-augmenting technology is *labor-saving* and thus capable of inducing investment coincident with net employment contraction.

To introduce labor-augmenting technology in a meaningful way, we must depart from the Cobb-Douglas production function used above. Instead, we specify that production is given by the CES function,

$$\left(ak^{\frac{\varphi-1}{\varphi}} + (1-a)(\xi n)^{\frac{\varphi-1}{\varphi}}\right)^{\frac{\varphi}{\varphi-1}},$$

where φ is the elasticity of substitution across the factors, ξ is labor-augmenting technical change, and a (partially) governs capital's share of output. This function displays constant returns, under which plant size is indeterminate. To have a notion of plant size, we assume the firm faces an isoelastic product demand schedule given by $\zeta p^{-\epsilon}$, where ζ is the demand shifter. It follows that the value of the firm becomes

$$\Pi(k_{-1}, n_{-1}, \xi, \zeta) = \max_{k,n} \left\{ \begin{array}{c} \zeta^{1/\epsilon}\left(ak^{\frac{\varphi-1}{\varphi}} + (1-a)(\xi n)^{\frac{\varphi-1}{\varphi}}\right)^{\frac{\varphi}{\varphi-1}\frac{\epsilon-1}{\epsilon}} - wn - \mathcal{C}_k(k, k_{-1}) - \mathcal{C}_n(n, n_{-1}) \\ +\mathcal{D}\int\int \Pi(k, n, \xi', \zeta') f_\xi(\xi'|\xi) f_\zeta(\zeta'|\zeta)\, d\xi d\zeta, \end{array} \right\},$$

where, for any ϱ, f_ϱ is the conditional p.d.f. of ϱ. It is useful to assume, as we did above with regard to the Hicks-neutral technology, that the demand shifter follows a geometric random walk. Then one may normalize the factors with respect to ζ, giving $\tilde{n} \equiv n/\zeta$ and $\tilde{k} \equiv k/\zeta$, and show that normalized problem, recast in terms of $\left(\tilde{k}, \tilde{n}, \xi, \zeta\right)$, is linearly homogeneous in ζ. Thus, the policy rules given in section 1.1 hold here, *conditional on* ξ. So there is an analogue to Figure 1; we must simply bear in mind that it is contingent on a given value of ξ.

To gain some intuition for the role of ξ, it may be helpful to revert briefly to the static problem. In that case, it is straightforward to derive the elasticity of employment with respect to labor-augmenting technical change. It is

$$(\epsilon - 1) - (\epsilon - \varphi)\frac{\lambda \xi^{1-\varphi}}{\lambda \xi^{1-\varphi} + 1 - a},$$

where λ is a nonlinear combination of a, φ, and factor prices. The first term, $\epsilon - 1$, captures the fact that an increase in ξ shifts the supply schedule to the right, reduces the price, and raises output. The second term appears because labor-augmenting technical change raises, on impact and all else equal, the labor input to the firm. This increases the marginal

and Ramey (2005) show that a labor-saving innovation can induce the empirically observed pattern in aggregate impulse responses: investment rises and employment declines. They confine the analysis to the limiting case of Leontif production. We will consider a middle ground where the elasticity of substitution is positive but less than one.

product of capital. If $\varphi < \epsilon$, this, in turn, induces the firm to substitute, to some extent, capital for labor. This latter effect will dominate the former if φ is relatively low. In this case, labor-augmenting technical change represents a labor-saving innovation: n declines in ξ. Intuitively, if φ is low, the two factors are more complementary in production, which implies that the immediate increase in capital demand after the arrival of ξ is very high.

This discussion suggests that, in the dynamic problem, an increase in ξ shifts the parallelogram in Figure 1 down and to the right if φ is sufficiently low. This is shown in Figure 7. An increase in ξ makes the firm more inclined to invest, so the capital adjustment barriers shift to the right. At the same time, if φ is sufficiently low, the firm is more inclined to reduce its size, so the employment barriers shift down. As a result, if the pair $\left(\tilde{k}, \tilde{n}\right)$ was at point A before the increase in ξ, it lies to the northwest of the parallelogram after the change in ξ. This corresponds to the area where it is optimal to both invest and to fire.[33]

[Simulation analysis to be completed]

4 Related literature

This paper is related to recent work by Eslava, Haltiwanger, Kugler, and Kugler (2010). These authors build off the idea that, at least under certain conditions, the firm's problem may be re-cast in terms of the gaps between the actual and "desired" levels of capital and employment. But rather than solve a fully specified structural model, they adopt a more reduced-form approach.[34] Their results do indicate certain interactions across the factor demands. For instance, the probability of expanding employment is increasing in the firm's (initial level of) capital, and the probability of investing is increasing in the firm's (initial level of) employment.[35]

[33] Dixit (1997, top of page 18) appears to have anticipated this point in some of his concluding remarks on the model.

[34] To elaborate, consider a (Brownian) model in which employment is the *lone* factor of production, and it is subject only to a *fixed* cost of adjusting. One may show that a firm follows an Ss policy such that, whenever it adjusts, it sets employment so that $n/\bar{n} = m$, where m is a constant (to be chosen optimally) and $\bar{n}$ is the static optimum. Hence, $n^* \equiv m\bar{n}$ is thought of as "desired" employment. It follows that, in this setting, one can calculate the logarithmic gap between actual and desired employment simply by using data on n and an estimate of the *static* policy. Eslava, et al (2010) apply this approach individually to each factor, estimating the gaps between actual and desired employment (n/n^*) and capital (k/k^*) using the static optima. They then estimate the movements in the one gap (say, n/n^*) associated with movements in the other (say, k/k^*). However, in a *multi*-factor model with *both* linear and fixed costs of control associated with each factor, the microeconomic foundation of this gap approach is not clear. This highlights the advantage of fully specifying and solving a structural model.

[35] The authors do not plot a graph like Figure 4, so we cannot comment on whether their data are consistent with the model of section 1.

We view our approach to the dual factor adjustment problem as complementary. Rather than assume a parametric policy function, we take a structural approach by specifying the precise adjustment frictions faced by the firm. The payoff is that we obtain, in the case of linear costs of adjustment, a clear testable implication about the qualitative behavior of capital and employment at the plant level. Though the presence of fixed costs complicates the analytics, we can still uncover reasonably robust predictions regarding the co-movement of the two factors via numerical analysis. As a result, we can perform a rigorous test of this specific class of adjustment cost models.

Another related paper is from Sakellaris (2000). This paper studies labor and capital adjustment in the U.S. Annual Survey of Manufacturers. It also finds that, in periods in which an investment spike occurs, a large decline in employment still appears to occur with the same frequency as a large rise in employment. This echoes our result. However, this result is not a focus of Sakellaris' (2000) study, and he does not relate it to the empirical predictions of a costly adjustment model, as we do above.[36]

5 References

Abel, Andrew and Janice Eberly (1996). "Optimal Investment with Costly Reversibility," *Review of Economic Studies*, Vol. 63, No. 4 (October), pp. 581-594.

Anderson, Patricia (1993). "Linear Adjustment Costs and Seasonal Labor Demand: Evidence from Retail Trade Firms," *Quarterly Journal of Economics*, Vol. 108, (November), pp. 1015-1042.

Barro, Robert J. and Robert G. King (1984). "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles." *Quarterly Journal of Economics*, Vol. 99, No. 4, (November), pp. 817- 839.

Bentolila Samuel and Guiseppe Bertola (1990). "Firing costs and labour demand: How bad is Eurosclerosis?", *Review of Economic Studies*, 57(3), July, pp. 381-402.

Bils, Mark and Peter Klenow (2004). "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, Vol. 112 (October), pp. 947-985.

Bloom, Nicholas (2009). "The impact of uncertainty shocks," *Econometrica*.

Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten and Stephen Terry (2011). "Really Uncertain Business Cycles," mimeo.

[36] Indeed, Sakellaris mentions the result only in a footnote in the working paper version (2000). The final published paper (Sakellaris (2004)) omits any discussion of the finding and its implications for dynamic factor demand models.

Chirinko, Robert (2008). "σ: The long and short of it," *Journal of Macroeconomics*, Vol. 30, No. 2 (June), pp. 671-686.

Cooper, Russell and John Haltiwanger (2006). "On the Nature of Capital Adjustment Costs", *Review of Economic Studies*, vol. 73, pp. 611-633.

Cooper, Russell, John Haltiwanger, and Laura Power (1999). "Machine Replacement and the Business Cycle: Lumps and Bumps," *American Economic Review.*

Cooper, Russell, John Haltiwanger, and Jonathan Willis (2007). "Search frictions: Matching aggregate and establishment observations," *Journal of Monetary Economics.*

Davis, Steven and John Haltiwanger (1992). "Gross job creation, gross job destruction, and employment reallocation," *Quarterly Journal of Economics*, pp. 819-864.

Dixit, Avinash (1997). "Investment and Employment Dynamics in the Short Run and the Long Run," *Oxford Economic Papers*, Vol. 49, No. 1 (January), pp. 1-20.

Doms, Mark and Timothy Dunne (1998). "Capital adjustment patterns in manufacturing plants," *Review of Economic Dynamics*, Vol. 1, pp. 409-429.

Eberly, Janice and Jan van Mieghem (1997). "Multi-factor Dynamic Investment Under Uncertainty," *Journal of Economic Theory*, Vol. 75, No. 2 (August), pp. 345-387.

Eslava, Marcela, John Haltiwanger, Adriana Kugler, and Maurice Kugler (2010). "Factor Adjustment After Deregulation: Panel Evidence from Colombia Plants," *Review of Economics and Statistics* (May).

Feldstein, Martin and Michael Rothschild (1974). "Towards an economic theory of replacement investment," *Econometrica*, Vol. 42, No. 3 (May), pp. 393-424.

Foster, Lucia, John Haltiwanger, and Chad Syverson (2008). "Reallocation, Firm Turnover and Efficiency: Selection on Productivity or Profitability," *American Economic Review.*

Francis, Neville and Valerie Ramey (2005). "Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited," *Journal of Monetary Economics*, Vol. 52, No. 8 (November), pp. 1379-1399.

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

Klenow, Peter and Benjamin Malin (2011). "Microeconomic Evidence on Price-Setting," *Handbook of Monetary Economics* 3A, B. Friedman and M. Woodford ed.: Elsevier, pp. 231-284.

Kumhof, Michael and Douglas Laxton (2009). "Chile's Structural Fiscal Surplus Rule: A Model-Based Evaluation," IMF WP.

Levinsohn, James (1999). "Employment Responses to International Liberalization in Chile," *Journal of International Economics*, Vol. 47 (April), pp. 321-344.

Li, Zhe (2009). "Lumpy Capital, Labor Market Search and Employment Dynamics over Business Cycles," mimeo.

Liu, Lili (1993). "Entry-exit, learning, and productivity change: Evidence from Chile," Journal of Development Economics, Vol. 42, No. 2 (December), pp. 217-242.

McCarthy, Jonathan and Egon Zakrajsek (2000). "Microeconomic Inventory Adjustment: Evidence from U.S. Firm-level Data," Finance and Economics Discussion Series 2000-24.

Mosser, Patricia (1990). "Trade Inventories and (S,s)," *Quarterly Journal of Economics*.

Nakamura, Emi and Jon Steinsson (2008). "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, Vol. 123, No.4, pp. 1415-1464.

Oh, Jiyoon (2011). "The cyclicality of firm size distribution and its effect on aggregate fluctuations," mimeo.

Petrin, Amil and Jagadeesh Sivadasan (forthcoming). "Estimating Lost Output from Allocative Inefficiency, with Application to Chile and Firing Costs," *Review of Economics and Statistics*.

Ramey, Valerie and Matthew Shapiro (2001). Displaced Capital: A Study of Aerospace Plant Closings," *Journal of Political Economy*.

Reiter, Michael, Tommy Sveen, and Lutz Weinke (2010). "Lumpy Investment and State-Dependent Pricing in General Equilibrium," mimeo.

Sakellaris, Plutarchos (2000). "Patterns of Plant Adjustment," mimeo.

Sakellaris, Plutarchos (2004). "Patterns of Plant Adjustment," *Journal of Monetary Economics* Vol. 51, pp. 425–450.

Figure 1: The optimal policy

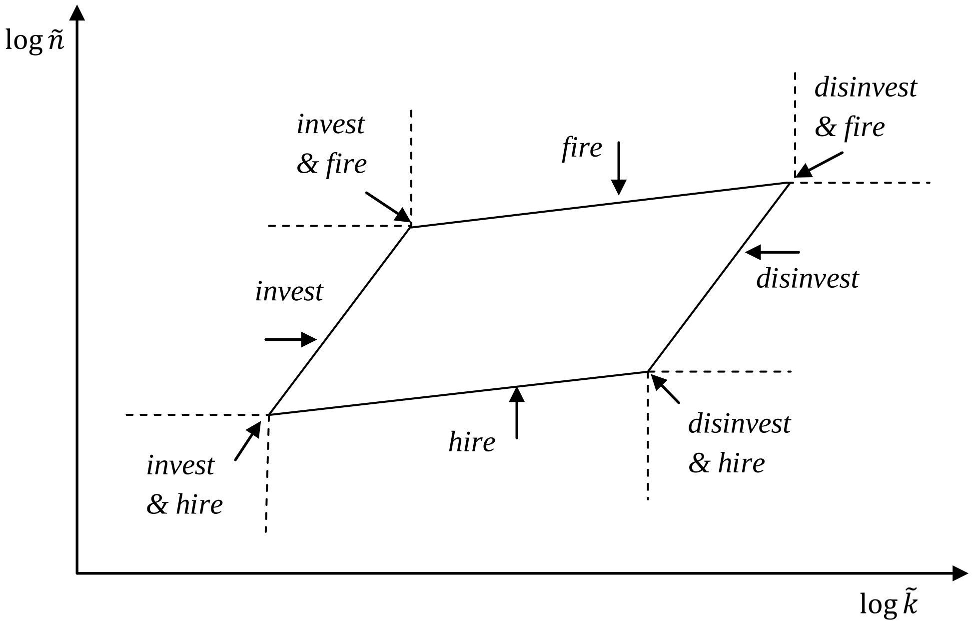

Figure 2: The joint dynamics of capital and employment

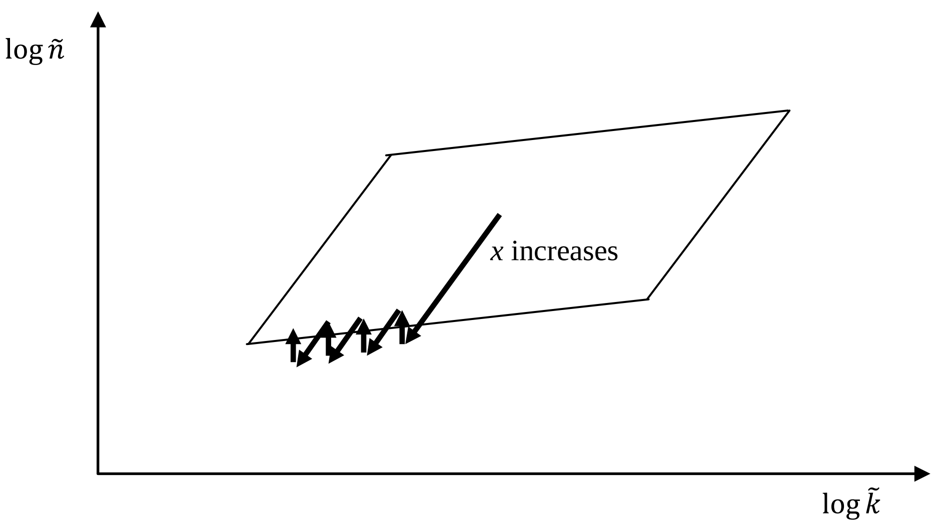

Figure 3: Policy rules and dynamics under fixed adjustment costs

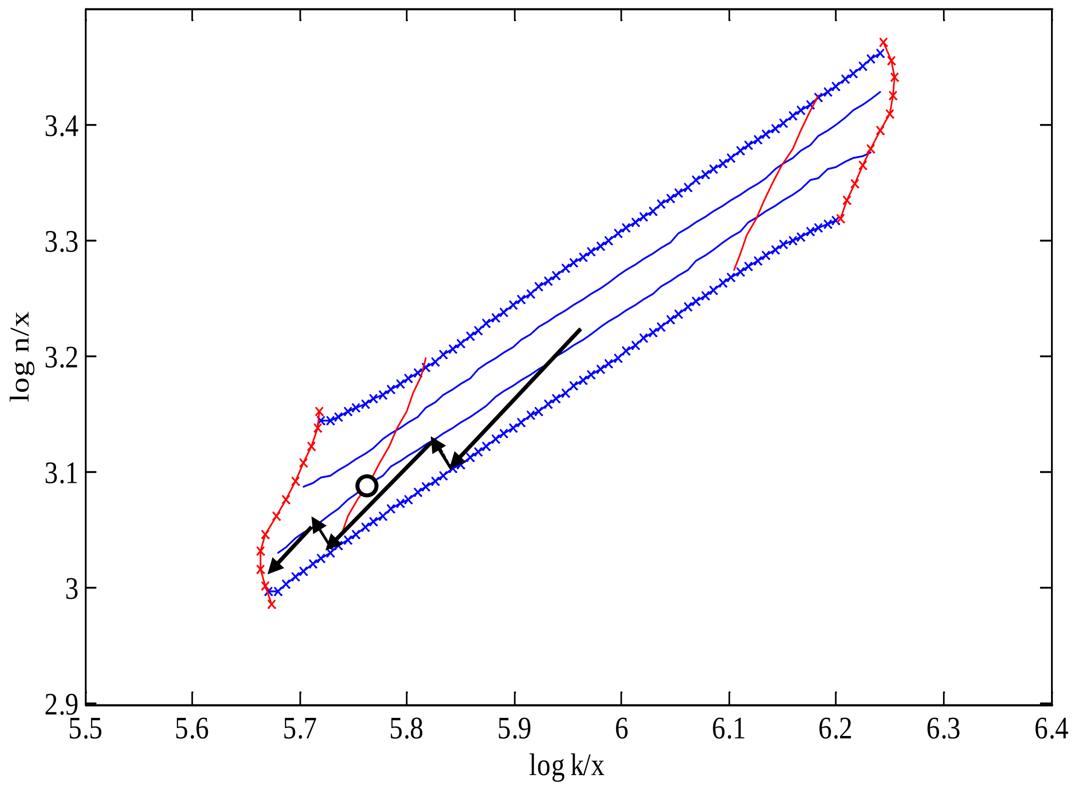

NOTE: The solid lines that make up the inner parallelogram show the reset policy rules: if a factor is adjusted, it is reset to a point along these schedules. The lines marked by "x", which make up the outer hexagon, show the triggers. If n/x or k/x reach one of the triggers, then the firm undertakes adjustment and moves the factor to its reset schedule. The arrows display the path of a pair $\left(\frac{n}{x}, \frac{k}{x}\right)$ as productivity, x, repeatedly increases. The circle marks the point to which the pair $\left(\frac{n}{x}, \frac{k}{x}\right)$ is adjusted when it reaches the southwest corner of the hexagon.

Figure 4: The empirical distribution of employment growth conditional on investment

A. Chile

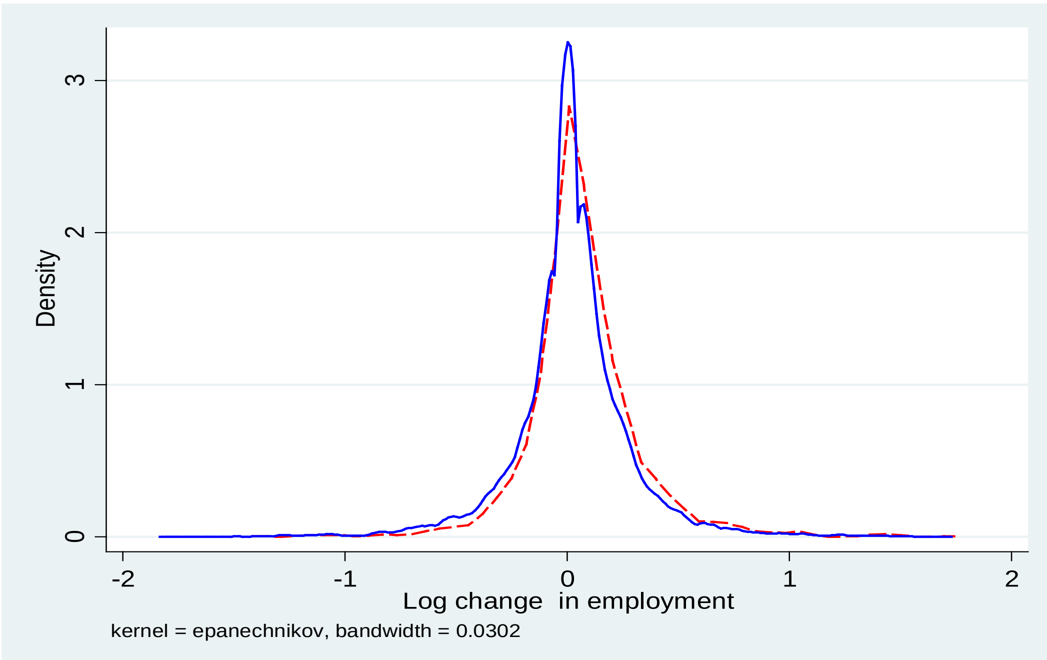

B. South Korea

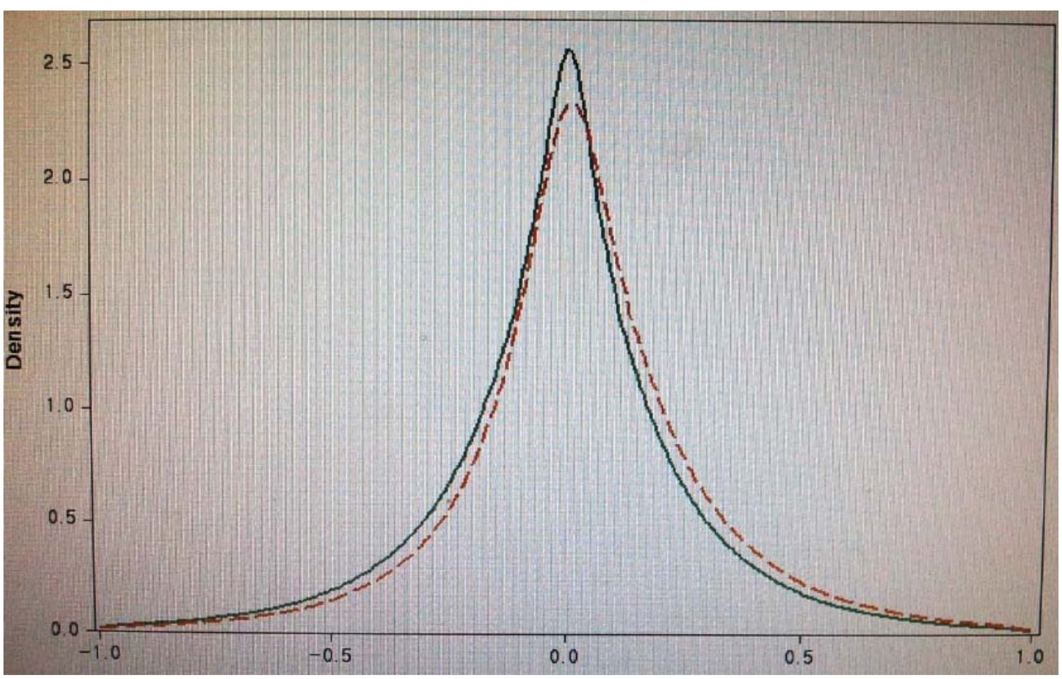

NOTE: In each panel, the solid line is the unconditional distribution of the log change in employment, and the dashed line is the distribution conditional on $\frac{I}{K_{-1}} > 10\%$.

Table 1: Baseline Calibration

Parameter	Meaning	Value	Reason
β	Elasticity of output wrt n	0.60	Labor share
α	Elasticity of output wrt k	0.25	Capital share
$\mathcal{D}$	Discount factor	0.98	Annual real interest rate = 8.4%
c_k^+	Purchase price of capital	1	Normalization
$-c_k^-$	Resale price of capital	0.96	Frequency of investment
c_n^+, c_n^-	Cost to hire & fire	28.8% of monthly wage	Frequency of net employment adj.
σ	Std. dev. of innovation to x	0.15	Avg. increase in $\log n$, given $I/K_{-1} > 10\%$

Table 2: Simulated and empirical moments, baseline model

Moment	**Model**	**Data** *Chile*	*Korea*
Prob($\Delta n = 0$):	0.1216	0.144	0.142
Prob$\left(\frac{I}{K_{-1}} = 0\right)$:	0.473	0.466	0.497
Std dev of $\Delta \log n$:	0.1828	0.264	0.31
Prob$\left(\frac{I}{K_{-1}} < 0 \mid \frac{I}{K_{-1}} \neq 0\right)$:	0.4944	0.034	0.0574
Share of est. contracting n qtly., given $\frac{I}{K_{-1}} > 0$:	0	n.a.	n.a.
Avg. decrease in $\log n$ qtly., given $\frac{I}{K_{-1}} > 0$:	0	n.a.	n.a.
Avg. increase in $\log n$ qtly., given $\frac{I}{K_{-1}} > 0$:	0.097	n.a.	n.a.
Share of est. contracting n yrly., given $\frac{I}{K_{-1}} > 0.1$:	0.0004	0.304	0.47
Avg. decrease in $\log n$ yrly., given $\frac{I}{K_{-1}} > 0.1$:	-0.0174	-0.164	-0.218
Avg. increase in $\log n$ yrly., given $\frac{I}{K_{-1}} > 0.1$:	0.282	0.207	0.244

NOTE: The top panel presents moments related to the annual distribution of investment and employment growth across plants. The moments (in this order) are the probability of not adjusting employment; the probability of not adjusting capital; the standard deviation of the distribution of the annual log change in employment; and the share of observations that involve disinvestment, conditional on adjusting capital. The second panel presents three moments related to the quarterly distribution of employment growth, conditional on positive investment. The third panel presents those same three moments, computed this time from the annual (yearly) distribution of employment growth.

Figure 5: The (baseline) model-implied and empirical distribution of employment growth conditional on investment

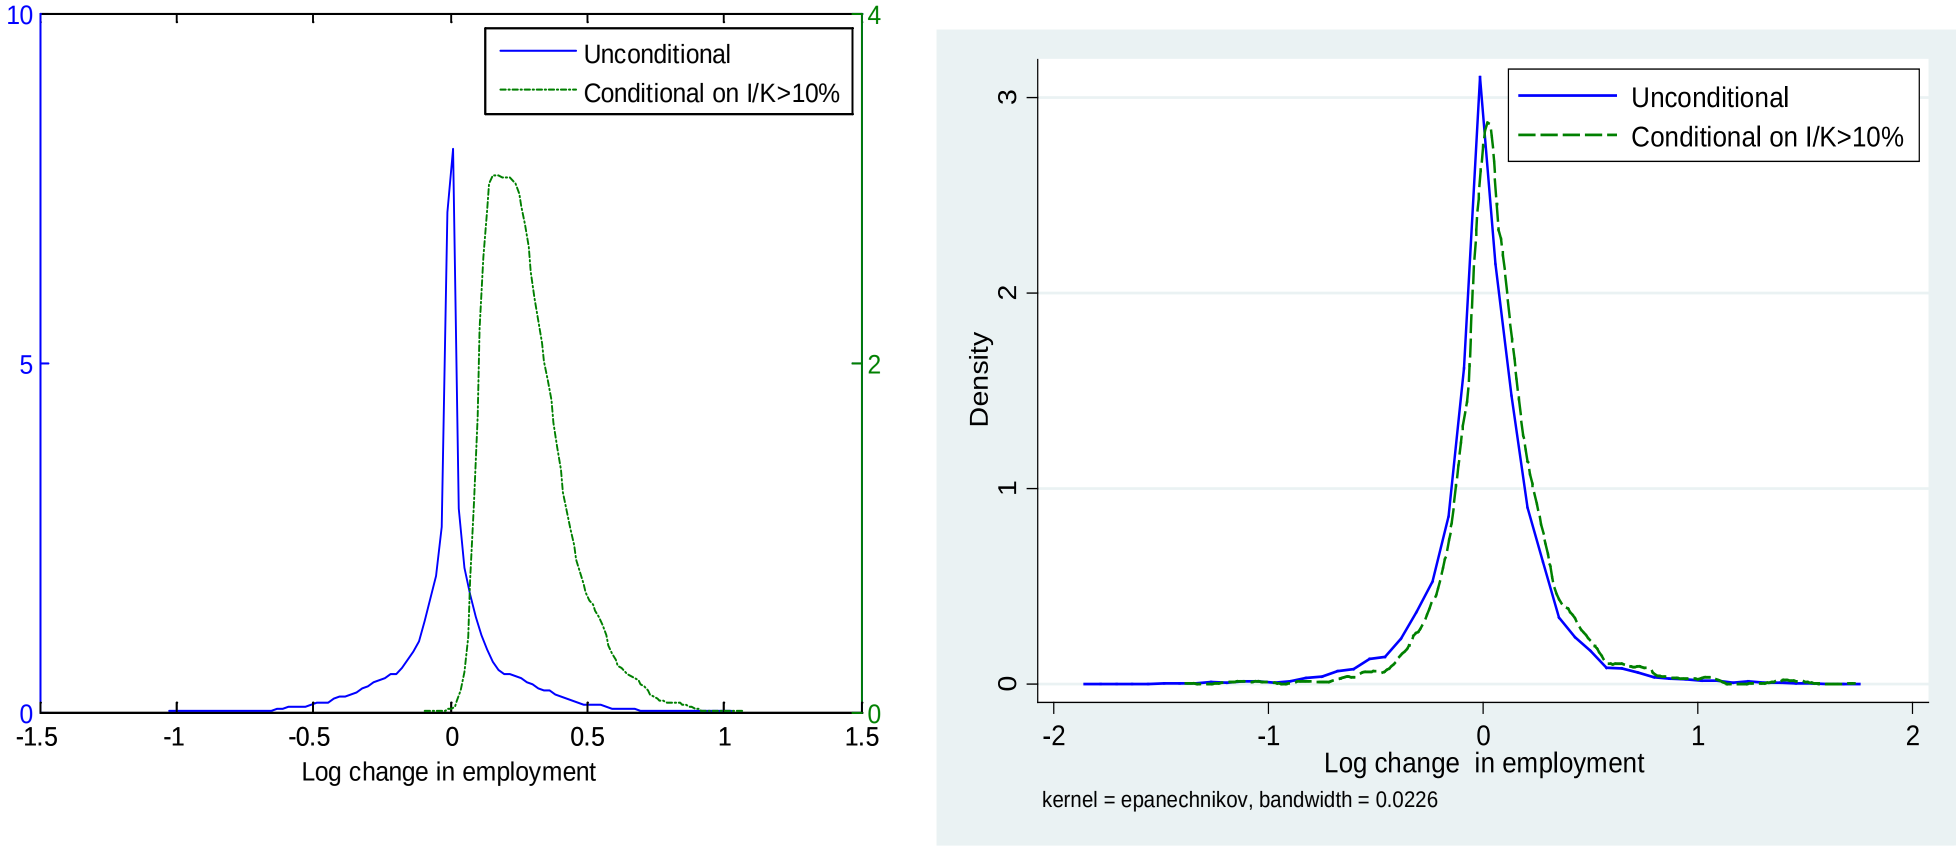

NOTE: The left panel shows the model-implied distribution of the annual log change in employment. The right panel shows the empirical distribution for Chile.

Table 3: Simulated and empirical moments, sensitivity analysis

	Model		**Data**	
	Constant	*Stochastic*		
Moment	$\delta_k > 0$	$\delta_k > 0$	*Chile*	*Korea*
Prob($\Delta n = 0$):	0.113	0.108	0.144	0.142
Prob$\left(\frac{I}{K_{-1}} = 0\right)$:	0.437	0.464	0.466	0.497
Std dev of $\Delta \log n$:	0.173	0.176	0.264	0.31
Prob$\left(\frac{I}{K_{-1}} < 0 \mid \frac{I}{K_{-1}} \neq 0\right)$:	0.0246	0.028	0.034	0.0574
Share of est. contracting n qtly., given $\frac{I}{K_{-1}} > 0$:	0.012	0.037	n.a.	n.a.
Avg. decrease in $\log n$ qtly., given $\frac{I}{K_{-1}} > 0$:	-0.003	-0.0286	n.a.	n.a.
Avg. increase in $\log n$ qtly., given $\frac{I}{K_{-1}} > 0$:	0.101	0.088	n.a.	n.a.
Share of est. contracting n yrly., given $\frac{I}{K_{-1}} > 0.1$:	0.056	0.137	0.304	0.47
Avg. decrease in $\log n$ yrly., given $\frac{I}{K_{-1}} > 0.1$:	-0.025	-0.063	-0.164	-0.218
Avg. increase in $\log n$ yrly., given $\frac{I}{K_{-1}} > 0.1$:	0.204	0.211	0.207	0.244

NOTE: In regards to the definitions of the moments, see the Note to Table 2. The column labeled "constant $\delta_k > 0$" presents simulation results for the model with a constant rate of depreciation. The depreciation rate is set to 2.5% per quarter. The column labeled "stochastic $\delta_k > 0$" presents simulation results for the model with stochastic depreciation. See text for details of the calibration.

Table 4: Simulated and empirical moments, sensitivity analysis cont.

	Model		**Data**	
Moment	$\delta_n > 0$	FC	*Chile*	*Korea*
Prob($\Delta n = 0$):	0.003	0.13	0.144	0.142
Prob$\left(\frac{I}{K_{-1}} = 0\right)$:	0.466	0.441	0.466	0.497
Std dev of $\Delta \log n$:	0.183	0.237	0.264	0.31
Prob$\left(\frac{I}{K_{-1}} < 0 \mid \frac{I}{K_{-1}} \neq 0\right)$:	0.55	0.494	0.034	0.0574
Share of est. contracting n qtly., given $\frac{I}{K_{-1}} > 0$:	0	0	n.a.	n.a.
Avg. decrease in $\log n$ qtly., given $\frac{I}{K_{-1}} > 0$:	0	0	n.a.	n.a.
Avg. increase in $\log n$ qtly., given $\frac{I}{K_{-1}} > 0$:	0.118	0.191	n.a.	n.a.
Share of est. contracting n yrly., given $\frac{I}{K_{-1}} > 0.1$:	0.0002	0.005	0.304	0.47
Avg. decrease in $\log n$ yrly., given $\frac{I}{K_{-1}} > 0.1$:	-0.008	-0.021	-0.164	-0.218
Avg. increase in $\log n$ yrly., given $\frac{I}{K_{-1}} > 0.1$:	0.286	0.276	0.207	0.244

NOTE: In regards to the definitions of the moments, see the Note to Table 2. The column labeled $\delta_n > 0$ presents simulation results for the model with attrition (and depreciation set to zero). The attrition rate is fixed at 2.5% per quarter. The column labeled "FC" presents simulation results for the model with fixed costs of capital and employment adjustment (and depreciation and attrition set to zero).

Figure 6: Dynamics under stochastic depreciation

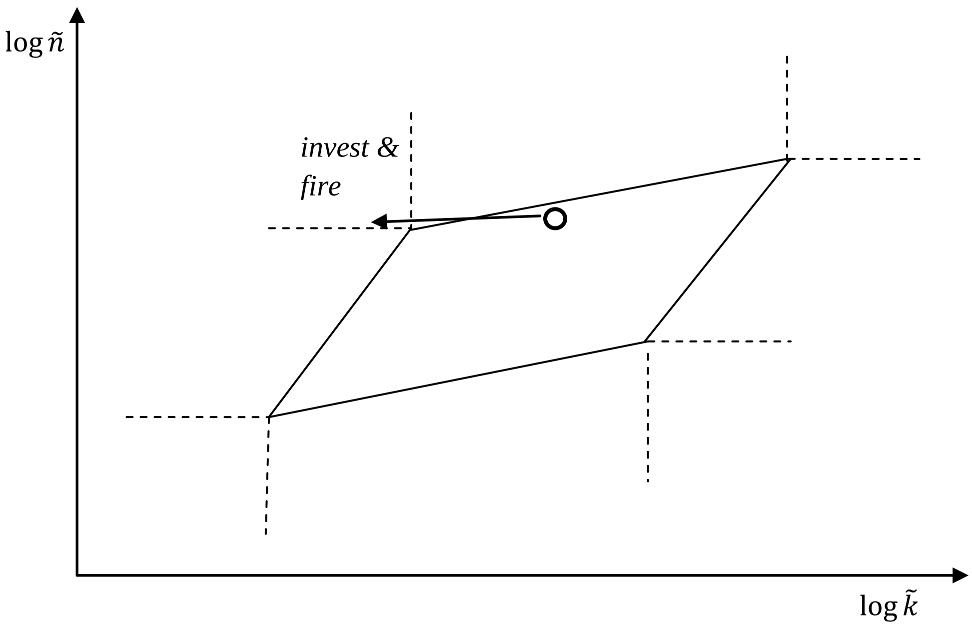

Figure 7: The reaction of the optimal policy to an increase in labor-saving innovations