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

William Hawkins, Yeshiva University

Ryan Michaels, University of Rochester

Jiyoon Oh, KDI

Preliminary and incomplete

April 1, 2013

Abstract

This paper uses plant-level data to document the joint adjustment of capital and employment. The data are analyzed through the lens of a model that integrates features from influential theories of costly capital and labor adjustment. Although the model is rich enough to account for several dimensions of individual factor dynamics, complementarity of the factors implies a strong restriction on their joint adjustment: investment ought to perfectly predict employment growth (Dixit, 1997). In contrast, 42 percent of gross capital accumulation in our data occurs at times when employment falls. The paper considers several extensions to the baseline model, but the key prediction is robust to, among other considerations, the introduction of delivery lags, alternative adjustment costs, and standard theories of factor-biased technical change. To engage the data, the paper finds that a more fundamental change in production technology might be needed, namely, a production function in which machinery directly replaces labor in certain tasks. The paper concludes by illustrating the macroeconomic implications of its findings.

*We are grateful to seminar participants at the University of Rochester and the Université de Montréal and to conference participants at the 2012 Comparative Analysis of Enterprise Data and the 2012 New York/Philadelphia Workshop on Quantitative Macroeconomics.

E-mail address for correspondence: ryan.michaels@rochester.edu.

Recent research in macroeconomics has emphasized models whose microeconomic environments are designed to be consistent with certain 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 fairly common 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 adjustment.[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 in section 1 the implications of a prototypical model of capital and employment adjustment that integrates features from influential theories of dynamic capital and labor demand. The model assumes piece-wise linear costs of adjustment on each factor, that is, the cost of adjusting is proportional to the size of the 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

[1] A number of papers have documented this fact on U.S. data. On investment, see Cooper and Haltiwanger (2006); on employment, 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). The literature on investment (see also Doms and Dunne, 1998) has stressed the skewness and kurtosis of the investment distribution more than inaction. But as we discuss below, inaction rates might be relatively low in the U.S. sample because it consists of relatively large establishments.

[2] Theoretical analyses of models with multiple frictions include Bloom (2009); Reiter, Sveen, and Weinke (2009); and Bloom, Floetotto, Jaimovich, Saporta-Eksten and Terry (2011). Empirical analyses include Eslava, Haltiwanger, Kugler, and Kugler (2010) and Sakellaris (2000, 2004). I discuss the latter papers in more detail below.

[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), Veracierto (2002), and Khan and Thomas (2011). Ramey and

model places testable restrictions on factor adjustment: the relatively less costly-to-adjust factor may be updated without changing the other, but if the relatively more costly-to-adjust factor is changed, complementarity across the factors in production implies that the less costly-to-adjust factor is *always* updated.

The data analysis of section 2 reveals that the model's prediction is at odds with plant-level behavior. We use annual establishment-level data on machinery investment and employment from the censuses of manufacturing in Chile and South Korea.[4] The empirical frequencies of adjustment indicate that capital is the more costly-to-adjust factor.[5] Hence, the model implies that investment perfectly predicts positive employment growth. Yet in the data, we find that between 29.5 percent (Chile) and 39.5 percent (Korea) of plants that undertake investment also reduce employment.

This coincidence of investment with declines in employment is *not* due to episodes in which one of the factors changes only slightly. For example, among plants which invest, the declines in employment at those which reduce their workforce are almost as large as the expansions in employment at those which increase their workforce. Moreover, if we cumulate all of the investment undertaken at times when employment falls, the total amounts to 42 percent of aggregate gross capital accumulation in our sample.

Our result does not contradict research which finds that, *on average*, plant-level employment growth is higher conditional on investment (Letterie, et al, 2004; Eslava, et al, 2010). We stress, however, that the class of theoretical models which often provides the conceptual underpinnings of earlier empirical analysis implies much stronger restrictions on the joint dynamics of capital and labor demand.[6] In this sense, we have sought to provide a more rigorous test of this class of models.

To assess the sensitivity of this result, we first vary the threshold for what constitutes positive net investment. To limit instances where zero investment is mis-reported as small positive investment, our baseline analysis conditions on an investment rate (annual investment divided by end-of-last-year capital) in excess of 10 percent. We also consider a threshold

Shapiro (2001) give direct evidence on the cost of reversibility, and Cooper and Haltiwanger (2006) estimate the degree of irreversibility in a structural model. Linear costs of adjusting employment have been considered in Bentolila and Bertola (1990), Anderson (1993), and Veracierto (2008).

[4] Throughout, we use "establishment" and "plant" interchangeably. The notion of a "firm" is very distinct. We discuss later the advantages of plant-level, as opposed to firm-level, data.

[5] Within the model of section 1, the relatively costly-to-adjust factor is, in essence, the factor whose Ss bands are wider. When we calibrate the model to replicate the adjustment frequencies of the individual factors observed in the data, the result is that the bands of the investment policy rule are further apart.

[6] Eslava, Haltiwanger, Kugler, and Kugler appear to have in mind a model where there are (at least) fixed costs of adjusting each factor. The addition of fixed costs complicates the analytics, but, as we show later, the quantitative implications are unaffected.

of 20 percent, which is used in the related literature to define an investment "spike" (Cooper and Haltiwanger, 2006). The results are largely unaffected.

We then investigate forms of aggregation bias. For instance, plants are arguably aggregates over at least somewhat heterogeneous production units. If one unit reduces employment substantially while another invests, the establishment as a whole is seen to reduce its labor demand even as it invests. It is difficult to address this issue conclusively. We do note, though, that if this aggregation bias were pronounced, the coincidence of positive investment and net separations would become negligible at smaller plants, where there is less scope for heterogeneous operations. However, the result holds for both small and large establishments.

A second source of aggregation bias is time aggregation. Specifically, the establishment may invest in one quarter, but contract employment significantly later in the year. We can only address this concern by simulation analysis. We calibrate the model at a quarterly frequency and select structural parameters to replicate the annual adjustment frequencies of individual factors. We ask whether the model replicates the relatively weak correlation between annual investment and employment adjustment and find that it does not.

Section 3 discusses a number of possible theories for the comovement between capital and labor. We conclude that many of them do not provide fully satisfactory accounts of the data. We mention a few here. First, the measured co-movement between machinery investment and employment growth weakens if there are lags in the delivery of capital goods and if plants record investment only when the equipment arrives. In this case, a plant may place an order now but by the time the new machinery is delivered, TFP (or product demand) has fallen to the point where desired employment is lower. However, the available evidence suggests a fairly modest average delivery lag of around 6 months. As a result, when we generalize the baseline model to allow for delivery lags, the coincidence between positive investment and negative employment growth in the model-generated data remains significantly below what we find in the (actual) data.[7]

Second, the coincidence of positive investment and negative employment growth suggests factor-biased technical change. To address this, we first modify the baseline model to include a constant-elasticity-of-substitution (CES) production function with labor-augmenting technology. Technical progress in this model stimulates investment[8] and, since it reduces the required labor input for a *given* level of output, depresses labor demand. However, unless the plant faces a very low price elasticity of demand for its product, the increase in output that

[7] Delivery lags in machinery are smaller than the time-to-build in the case of nonresidential structures. This is one reason we focus on machinery investment. For a discussion of time-to-build, see Edge (2000).

[8] This is so if machinery and effective labor input are gross complements. Raval (2012) presents evidence for this.

accompanies the improvement in technology typically leads, on net, to higher labor demand. For instance, even if product demand is so inelastic as to imply a mark-up of 70 percent, an increase in labor-augmenting TFP still triggers an increase in labor demand. The same argument applies, as we show, in the case of capital-augmenting technical change.[9]

The paper next considers the form of technical progress studied in Krusell, et al (2000). These authors analyze a model of investment-specific technical change in which the technology embodied in new capital is a gross complement to skilled labor and a gross substitute to unskilled labor. This suggests the decline in plant-wide employment may reflect a shift against unskilled labor. To investigate this, we can distinguish between production and non-production workers; the literature has typically treated non-production status as a proxy for high skill (see Berman, Bound, and Machins, 1998). In contrast to the model's prediction, we find that there are nearly as many episodes of positive investment and declines in non-production workers as there are episodes of positive investment and declines in production workers.

Lastly, we study a form of technical change that represents a more marked deviation from the older literature. We adapt a recent model of Acemoglu (2010) in which labor can be directly replaced by machinery in tasks. In this model, an increase in capital-augmenting technology raises the effective supply of capital available to the firm, which, in turn, lowers its (internal) shadow price. This motivates the firm to deploy machinery to take over more tasks within the plant. This directly reduces, in absolute terms, the contribution of labor to production. In other words, capital-augmenting technical change in a model of tasks represents an adverse shift in labor productivity. To the best of our knowledge, this is the first time this model has been applied to account for plant-level factor adjustment dynamics.

The task-based model has much more success engaging the plant-level data. In particular, it very nearly replicates both the incidence of positive investment and negative employment growth and the size of the contractions in employment in these episodes. The challenge faced by the model is that, in order to target these moments, it typically – and, given its structure, quite naturally – implies a very high elasticity of substitution. This may prove difficult to reconcile with the evidence from U.S. plants in Oberfield and Raval (2012). These authors do not find that labor income shares react strongly enough to shifts in the relative price

[9] The assumption that the plant has some monopoly power, and so faces a downward-sloped demand schedule, is one way to introduce decreasing returns into the plant's revenue function. More generally, then, one must assume a significant degree of decreasing returns in order to generate the empirical coincidence between positive investment and negative employment growth. For instance, an elasticity of demand of 2.5 (used in simulations in section 3) implies a degree of decreasing returns of 0.6 (that is, a doubling of capital and labor leads to a $2^{0.6} \cong 1.52-$fold increase in revenue).

of labor (the wage rate relative to the capital rental price) for capital and labor to be so substitutable. But at the same time, their estimates are very hard to square with our results on factor adjustment. Clearly, further work in this area is needed.

Before we leave this topic of tasks, it is important to stress that, although the baseline model might be "fixed" by changing the production technology along these lines, the cost-of-adjustment framework has been critical to the analysis. Without it, we would have no way to engage the data given the prevalence of inaction in capital and labor. The cost-of-adjustment framework provides a means to develop clear testable implications when inaction is optimal and thus more fully unlocks the potential for plant-level data on factor adjustment to inform us as to the environment in which firms operate.

The paper's analysis concludes in section 4 with a brief discussion of the macroeconomic implications of our results. In recent years, there has been interest, among academics and policy advocates, in legislation that attempts to accelerate capital demand, such as investment subsidies. The popular interest in investment subsidies in particular tends to stem, though, from the belief that they will raise *labor* demand. Our plant-level data, however, has pushed us in the direction of a model of tasks in which labor demand can react negatively to reductions in the cost of capital. To investigate this further, we consider a simple extension of this model to allow for aggregate investment tax incentive. We provide an illustrative simulation that confirms that, in the model of tasks, an investment subsidy has depressive effects on aggregate labor demand.

Although we have data from just two countries, we believe the results of our empirical and theoretical analysis are relevant to industrialized economies more broadly. We say this, in part, because our empirical results echo findings from a few other studies. In Sakellaris' (2000) analysis of the U.S. Annual Survey of Manufacturers, for instance, he finds that in years of investment spikes (when the net investment rate exceeds 20 percent), large declines in employment (declines in excess of 10 percent) occur with the same frequency as large increases in employment.[10] Polder and Verick (2004) studied German and Dutch data and also observed that employment often declines when investment is positive.[11] Our contribution, relative to these papers, is to relate the results to the specific predictions of the baseline factor demand model of section 1 and to explore extensions of that model that may reconcile

[10] Sakellaris mentions the result only in the working paper version (2000) of his paper. The final published paper (Sakellaris (2004)) omits any discussion of the finding.

[11] We should note that in Polder and Verick (2004), the inaction rate on investment is lower than the inaction rate on employment – labor is the harder-to-adjust factor. In this case, the canonical model suggests that, conditional on positive investment, plants ought to always hire. This prediction is clearly violated in their data.

theory and evidence.

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 the implications of time aggregation. Section 3 then takes up a number of possible extensions to the baseline model that may rationalize the empirical result. Section 4 analyzes the aggregate implications of an investment subsidy in a model in which increases in investment can coincide with contractions in labor demand. Section 5 concludes.

1 The firm's problem

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 employment, respectively, are assumed to be[12]

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

The cost (c_n^+) of expanding employment is often interpreted as the price of recruiting and/or training. The cost (c_n^-) of contracting employment may represent a statutory layoff cost.[13] Capital decisions can 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 that implies costly reversibility, i.e., $c_k^+ > -c_k^-$.

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

[13] In South Korea's two-tier labor market, the termination of a *permanent* worker does not entail a layoff cost per se – dismissals are regulated more directly. Labor law directs managers to "make every effort to avoid dismissal" of permanent workers and, if the plant is unionized, to engage in "sincere consultation" with the workers' representatives (Grubb, Lee and Tergeist, 2007). *Temporary* workers include those on fixed-term employment contracts. A fixed-term contract lasts for up to two years, but there are few restrictions on their renewal. The dismissal of fixed-term employees is largely unregulated. To retain tractability, we (and much of the literature) do not attempt to model labor market institutions in detail. Permanent and temporary workers are melded into a representative worker, and the adjustment frictions are assumed to act, in effect, as a simple tax on adjustments.

For now, we omit fixed costs of adjusting from (1), but we consider the effect of these later.

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

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

where x is plant-specific productivity, w is the wage rate and ϱ is the discount factor. We assume revenue is given by $x^{1-\alpha-\beta}k^{\alpha}n^{\beta}$. The Cobb-Douglas form is particularly tractable, but the analysis carries through as long as (k, n, x) are (q-) complements; 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 depreciation and attrition only to economize on notation and simplify (slightly) the presentation of the dynamics below. They do not affect the predictions of the model, and will be included in the quantitative assessment in the next section.[14]

Figure 1 summarizes the optimal policy and is taken from Dixit (1997).[15] The assumption of a random walk ensures that the problem is linearly homogeneous in $(k, k_{-1}, n, n_{-1}, x)$. As such, it admits a normalization.[16] 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.[17] 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},

[14] The only source of variation in (2) is idiosyncratic TFP shifts. Unless otherwise noted, we abstract from aggregate uncertainty. This is consistent with our current focus on the cross section rather than aggregate fluctuations.

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

[16] We have numerically investigated the behavior of the model when productivity is stationary. Persistent but stationary shocks do not overturn the main implication of the analysis.

[17] The form of the optimal policy follows from the concavity, supermodularity, and linear homogeneity of the value function. See Dixit (1997) and Eberly and van Mieghem (1997) for details. Our discussion in this section is based on their analysis.

is so low as to make firing optimal. That is, if we let

$$\tilde{\Pi}\left(k_{-1},n_{-1},x\right)\equiv x^{1-\alpha-\beta}k_{-1}^{\alpha}n_{-1}^{\beta}-wn_{-1}+\varrho\int\Pi\left(k_{-1},n_{-1},x'\right)\mathrm{d}G\left(x'|x\right),$$

then $\tilde{\Pi}_2\left(k_{-1},n_{-1},x\right)<-c_n^-$. At this point, the firm reduces employment to the point where n satisfies the first-order condition, $\tilde{\Pi}_2\left(k_{-1},n,x\right)=-c_n^-$. 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., $\tilde{\Pi}_2\left(k,n_{-1},x\right)>c_n^+$). 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 $\log\tilde{n}$ if its capital stock is larger.[18]

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.[19] 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 $\tilde{\Pi}_2\left(k_{-1},n,x\right)=c_n^+$.

If x continues to rise, the firm repeatedly hires in observance of its first-order condition.

[18]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 $\tilde{\Pi}_k\left(k,n,x\right)=c_k^+$ and $\tilde{\Pi}_n\left(k,n,x\right)=c_n^+$. 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.

[19]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.

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.*

Now suppose that x begins to decline. As a result, $\log \tilde{n}$ and $\log \tilde{k}$ reverse course and travel northeast through the parallelogram. The firing barrier is reached first. This is because the capital adjustment barriers are relatively far apart in the sense that the 45^0 degree line extending from the southwest corner crosses the firing barrier before reaching the disinvestment barrier. Along the firing barrier, the firm sets n to satisfy the first-order condition, $\tilde{\Pi}_2 (k_{-1}, n, x) = -c_n^-$. When the pair $\left(\log \tilde{k}, \log \tilde{n}\right)$ later reaches the disinvestment barrier, complementarity implies that both factors are reduced. If productivity begins to improve, the pair $\left(\log \tilde{k}, \log \tilde{n}\right)$ will again travel southwest toward the hiring barrier, and the process repeats.[20]

This argument has essentially traced the ergodic set of $\left(\log \tilde{k}, \log \tilde{n}\right)$ induced by the model. This is shown in the shaded region of the figure. This set stretches from the 45^0 ray from the southwest corner to the 45^0 ray that extends from the northeast corner of the parallelogram. The shape of the policy rules (and the structure of the stochastic process, x) implies that, any particle $\left(\log \tilde{k}, \log \tilde{n}\right)$ will eventually enter this space. Moreover, as our discussion has shown, the particle, once inside, will never leave.

2 Evaluation of the baseline model

2.1 The basic properties of the data

To assess the implications of the model, we use two sources of plant-level data. The first is the Korean Annual Manufacturing Survey, for which we have data from 1990-2006. The second is the Chilean Manufacturing Census, for which we have annual data from 1979-96. Both surveys cover all manufacturing establishments with at least 10 workers and include

[20] There is another, somewhat more subtle point to note, namely, the investment barrier has slope greater than one. Therefore, as $\left(\log \tilde{k}, \log \tilde{n}\right)$ travels northeast along the 45^0 line, it does not intersect the investment barrier before reaching the firing barrier. As Dixit (1997) and Eberly and van Meighiem (1997) observe, this property is implied by the homogeneity of the value function. Since capital satisfies the first-order condition $\tilde{\Pi}_1 (k, n, x) = c_k^+$ along the investment barrier, it follows that a perturbation to labor of $\mathrm{d} \log n$ yields a log change in capital demand, $\frac{\mathrm{d} \log k}{\mathrm{d} \log n} = \frac{\tilde{\Pi}_{12} n}{-\tilde{\Pi}_{11} k} > 0$. That $\tilde{\Pi}_1$ is homogeneous of degree zero restricts the size of the cross-partial, $\tilde{\Pi}_{12}$. Specifically, by Euler's Theorem, $\tilde{\Pi}_{12} n = -k\tilde{\Pi}_{11} - x\tilde{\Pi}_{13} < -k\tilde{\Pi}_{11}$, where the second inequality follows by supermodularity (e.g., $\tilde{\Pi}_{13} > 0$).

observations on the size of the plant's workforce and investment.[21] In what follows, we focus on machinery investment specifically, since that category is most consistent with the type of capital envisioned in the baseline model.

Though we have analyzed both datasets, we regard the South Korean survey as our principal source. The reason is that the measurement of employment in the Korean data is better suited to our purposes. Since investment is measured in both datasets as the cumulation of all machinery purchases over a calendar year, the analogous measure of employment growth is the change in employment between the end of the prior year and the end of the current year. However, employment in the Chilean Census is typically reported as an annual average. There are several consecutive years in the 1990s in which employment is reported as of the end of the (calendar) year, and our use of Chilean data must be restricted to this subsample. This is, in part, why the South Korean panel is so much larger. There are about 508,000 plant-year observations in the Korean survey between 1990-2006, which is an order of magnitude larger than than the (usable) Chilean panel.

Accordingly, our discussion in the main text focuses on the Korean data. Moreover, the results from Korean data will also serve as the targets of our calibration when we study the baseline model quantitatively. Appendix A catalogues the results from Chilean data. As noted in the introduction, the findings from Chilean data are largely similar to what we report off the Korean sample.

To take the model to data, we recall that the model's prediction pertains to employment growth conditional on *any* positive investment. However, measurement error in investment may mean that some of the smaller reported investments are in fact zeros. This is problematic insofar as negative employment growth conditional on *zero* investment is quite consistent with the baseline model. Hence, this form of measurement error could lead us to wrongly reject the model. Therefore, we wish to condition on investment rates in excess of 10 percent to guard against this. The investment rate is defined as i/k_{-1}, where i is real investment and k_{-1} is the end of the prior year real machinery stock.

The construction of a series for i/k_{-1} is done in the usual way via the perpetual inventory method. The capital stock in the first year of a plant's life in the sample is computed by deflating the book value of machinery by an equipment price index. Capital in the succeeding years is calculated using the law of motion, $k = (1 - \delta) k_{-1} + i$, assuming a depreciation rate of $\delta = 0.1$. Real investment, i, is obtained by deflating investment expenditure by the equipment price index. Note that each dataset includes information on gross purchases and

[21] After 2006, plants with $5 - 10$ workers were included in the Korean Census. For comparability across time, we use data only through 2006.

sales; investment expenditure is the difference of the two.[22]

Our exploration of the data begins with a characterization of the distributions of net investment and employment growth. This is shown in Table 1. Two features of the data are significant. First, judging by the adjustment frequencies of the individual factors, capital appears to be the more costly one to adjust. Investment is reported to be zero in South Korea manufacturing in 47 percent of all plant-year observations, whereas employment growth is zero in 13.5 percent of plant-year observations. (As shown in Appendix A, the telative frequency of investment is similar in Chile.) We will see that, when we calibrate the baseline model, these adjustment frequencies do in fact map to a relatively wide Ss band with respect to capital, which in turn makes capital the hard-to-adjust factor in the baseline model.[23]

The pervasiveness of inaction with respect to investment contrasts with estimates from other datasets. Cooper and Haltiwanger (2006) report that less than 10 percent of plant-year observations in their U.S. data show zero net investment. Letterie, Pfann, and Polder (2004) also report a small inaction rate in Dutch data. There are two points to bear in mind here.

It is important to bear in mind that the results based on U.S. and Dutch data are derived from a balanced sample of relatively large plants. For instance, Cooper and Haltiwanger's panel is related to that used in Caballero, Engel, and Haltiwanger (1997), and mean employment in the latter's sample was nearly 600.[24] Since large plants tend to be aggregations over at least somewhat heterogeneous production units, we anticipate that inaction rates are lower in these establishments. At plants with more than 100 workers, the share of plant-year observations in Korean data that involve no investment is 10.6 percent, and the share with zero employment change is 4 percent. These estimates are in the neighborhood of the numbers reported in the related papers. Note that, for large plants, investment is still the relatively less frequently adjusted factor.[25]

For most of the paper, we use our full sample. But, in an attempt to strike a compromise between our approach and that of others, we have re-run the main analysis – the distribution

[22] We use the price index for equipment purchases in the manufacturing sector. The price index is available from the OECD STAN database.

[23] To the degree that employment adjustment is less costly than capital adjustment in Korea, it is likely due to the fixed-term contract (see footnote 13 for more on this contract). However, our data do not identify fixed-term workers, so we are unable to investigate this more deeply.

[24] Letterie, Pfann, and Polder (2004) also use a balanced sample, so their sample is likely to consist of relatively large units. In related work by Polder and Verick (2004) on Dutch and German data, mean plant size is three (Dutch) and 9.5 (German) times that in Korean data.

[25] Of course, the adjustments at these larger plants contribute greatly to the change in aggregate employment. However, the approach in the literature has been to study the problem of a production unit, and the dynamics among small and mid-sized plants can inform our assessment of this model. To the extent there are heterogeneous units within a single plant, this aggregation problem ought to be addressed explicitly, but there is no consensus in the literature in this regard.

of employment growth conditional on positive investment – with a sample that excludes the first and last years of any plant's lifespan if that establishment enters or exits in our sample. The results hardly change. This should help address concerns that entry and exit drive the difference between the empirical moments and the model's predictions.

The second feature of the data that is worth notice has to do with the re-sale of machinery. Our data allows us to measure the gross sales of machinery by manufacturers. The frequency of re-sale in Korea (6.3 percent of sample observations) is roughly in line with that in U.S. (Cooper and Haltiwanger, 2006). More intriguingly, if a plant sells machinery, it is very likely to purchase equipment in the same year. The probability of an investment purchase conditional on positive sales in *the same year* is 74.6 percent. To the best of our knowledge, this fact has not been documented in the recent literature on factor adjustment. It is a fact that is at odds with the baseline model; the plant would never simultaneously sell and purchase the same machinery.[26] We suspect that this result suggests a kind of vintage-capital model in which a plant scraps or sells one machine and upgrades to a new vitnage (see Cooper, Haltiwanger, and Power (1999)). This fact is not the focus of this paper, but we return to the issue below since it might bear on the principal moment of interest, namely, the change in employment conditional on positive investment.

2.2 The principal result and its robustness

Given the relative frequency of investment, the baseline model predicts that investment (or disinvestment) is always accompanied by employment adjustment. The left panel of Figure 1 shows that it is not (we will discuss the right panel shortly). The figure shows the unconditional distribution of the net change in log employment across calendar years and the distribution conditional on a plant-level investment rate (i/k_{-1}) greater than 10 percent. The distribution of net employment growth, 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 that investment should perfectly predict an expansion in labor demand; employment growth should lie everywhere to the right of zero.[27]

[26] Although it is possible in the baseline model that a plant may purchase machinery in one month after a positive TFP realization but sell later in the year when a new productivity level is revealed, our work below suggests that this is unlikely given the size of the Ss bands on investment.

[27] Results in this section are based on episodes of (sufficiently) positive *net* investment. But since a plant may sell and purchase machinery in the same year, there will be more instances of positive gross investment than net investment. However, because re-sale is relatively infrequent, there are not that many instances in which gross investment exceeds 10 percent when net investment does not. Therefore, when we condition instead on gross investment rates greater than 10 percent, our results are largely unaffected.

To summarize the features of Figure 1, there are two moments that are particularly helpful and are reported in Table 1. First, among those plant-year observations that involve an investment rate greater than 10 percent, we compute the share in which employment declines. Second, we calculate the average contraction among the set of plants that both invest and "fire". In Korea, the former is 39.5 percent, and the latter is 19 percent (the median is 12.9 percent). To put the latter number in context, consider the average expansion in employment among plants that both invest and hire. This is 21.4 percent (the median is 15.4 percent). Thus, it is not the case that the baseline model is violated because small declines in employment coincide with positive investment – the declines are nearly as large as the increases in employment at expanding plants.

We now consider the robustness of these results. We are particularly concerned about aggregation bias. One concern that is relatively straightforward to address has to do industrial composition. In other words, perhaps there are (fixed) features of certain industries that affect the co-movement of capital and labor. We do not find any evidence of this. In Korea, for instance, we re-run the analysis by industry and find that the probability of negative employment growth conditional on positive investment is contained in a narrow range from 37 to 46 percent (see Table 2).

Second, the result may be partly due to aggregation over heterogeneous production units within large plants. 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.

This issue is difficult to address conclusively. We do note, however, that the scope for heterogeneous operations is presumably limited at smaller plants. Thus, within-plant aggregation were largely responsible for the result, the probability that employment declines conditional on positive investment should vanish at smaller plants. We do not find this. The bottom panel of Table 2 presents results by size class. The probability of employment contraction, conditional on positive investment, is lower at smaller plants: in Korea, for instance, it is 35.8 percent at plants with $10 - 24$ workers and 47.6 percent at plants with more than 100 workers. But the frequency with which employment falls in years of positive investment still appears economically significant.[28]

Third, the equipment stocks of plants are aggregates over heterogeneous machines. This

[28] 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, which includes only (relatively large) U.S. publicly traded corporations.

means it is possible that episodes of negative employment grwth and positive investment involve, predominantly, "small" investments, such as the replacement of hand tools. We have tried to minimize this possibility insofar as we condition on investment rates in excess of 10 percent. We have verified that the results are largely unaffected if we condition on investment rates greater than 20 percent, which is the threshold used to define investment spikes in Cooper and Haltiwanger (2006). The probability of negative employment growth, conditional on this 20 percent threshold, is 37.6 percent (it is 39.5 percent when the 10 percent threshold is used). The fact is that the investment undertaken in years of negative employment growth accounts for a significant fraction of total capital accumulation in our sample. When we cumulate investment in these episodes, it amounts to 42 share of all investment done.

Fourth, it is possible that the result reflects time aggregation. Establishments may invest and hire in one quarter 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, since we do not have higher-frequency data. In leiu of that, we can address this concern only via simulation analysis. We now detail our approach.

The objective is to simulate quarterly plant-level decisions from the model of section 1 and determine whether time-aggregating these observations to an annual frequency yields the joint adjustment dynamics observed in the annual data. There are five parameters introduced in section 1 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. In addition, since we want to allow for depreciation in the simulation, we must select the rate of decay, δ_k.

These structural parameters are chosen as follows. The price of the investment good, c_k^+, is normalized to one. The resale price, $-c_k^-$, is calibrated to target the frequency of investment. The idea 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, we impose symmetry on the costs of adjusting labor, $c_n^+ = c_n^-$ and choose this cost to target the frequency of net employment adjustment.[29] The variance of the innovation to productivity is used to target the average expansion in (log) employment among plants which increase employment and undertake investment. Given this, we ask whether the model can also generate the average *contraction* in employment

[29] The assumption $c_n^+ = c_n^- \equiv c_n$ may seem at odds with the sense that layoff restrictions are nontrivial in Korea (see footnote 13). However, without data on the gross flows of workers into and out of the plant, we can only target the distribution of net employment growth, and this is roughly symmetric about zero. The symmetry between c_n^+ and c_n^- allows us to replicate this.

among those plants which decrease employment and invest. Lastly, we fix $\delta_k = 0.02$.[30]

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 3. The discount factor, $\varrho = 0.9875$, is set to be consistent with the average real interest rate in South Korea over the sample for which we have data. The elasticity of output with respect to labor is assumed to be $\beta = 0.50$, which is consistent with the labor share of value-added in South Korean manufacturing.[31] To ensure a well-defined notion of plant size, the coefficient, α, attached to capital must then be set below $1 - \beta$. We fix $\alpha = 0.40$, so α/β is broadly consistent with the ratio of capital income to labor income in the sectoral accounts.[32] Note that for the moment, worker attrition is set to zero; we return to this matter below.[33]

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. (Appendix B reports the exact normalized model used in the simulations.)

Table 4 summarizes the simulation results. The table reports a few sets of statistics, each computed off the simulated panel. The top panel allows one to gauge how well the model matches the marginal distributions of each factor change. We report, for instance, the inaction rates with respect to employment and investment. (More precisely, these are the shares of establishment-year observations in the simulated panel for which there is no net change in employment and no investment, respectively.) These were moments targeted in

[30] All else equal, depreciation reduces the frequency of resale. If δ_k is high, the cost of reversing disinvestment (i.e., of having to pay a price for new machinery that exceeds what one can now earn by selling it) induces firms to shed capital exclusively via depreciation. Our choice of $\delta_k = 0.02$ implies too few sales in the baseline model, but we leave δ_k where it is since typical estimates of depreciation are not less than 2 percent.

[31] The real interest rate is calculated from OECD data on the short-term nominal interest and (realized) CPI inflation. Labor share can be computed directly off our census data; it is also publicly available through the EU KLEMS dataset. Note that in what follows, if we refer to sectoral accounts data, we mean EU KLEMS.

[32] To be precise, in the sectoral accounts, labor share's of value-added is measured to be 0.56 and capital's share is therefore 0.44. The difference between α, β, and these estimates implicitly represents the returns to fixed factors that, we assume, are imputed to capital owners and labor in the sectoral accounts. These may include the entrepreneurial capital of plant owners and/or managers and the value of land. The presence of of fixed factors implies a degree of decreasing returns that allows for a well-defined plant size.

[33] The use of value-added to compute labor share implicitly assumes a "true" production function of gross output of the form $x^{1-\bar{\alpha}-\bar{\beta}-\psi}k^{\bar{\alpha}}n^{\bar{\beta}}m^{\psi}$, where m is materials. If materials are costless to adjust and if the materials price is fixed, then one may concentrate out m and obtain $y = x^{\frac{1-\bar{\alpha}-\bar{\beta}-\psi}{1-\psi}} k^{\frac{\bar{\alpha}}{1-\psi}} n^{\frac{\bar{\beta}}{1-\psi}}$. The exponent attached to labor, $\beta \equiv \frac{\bar{\beta}}{1-\psi}$, is now interpretable as labor's share in value-added.

the calibration, so the model (nearly) reproduces the empirical estimates. In the top panel, we also report the unconditional standard deviation of the year-over-year change in log employment. This moment was not targeted, and the model does understate this measure of dispersion. An increase in σ would ameliorate this, but the presence of larger shocks would also imply larger average employment changes conditional on positive investment, which was targeted in the calibration. Note that the model does very nearly replicate this latter moment.

The next few rows summarize the joint adjustment of capital and employment at an annual frequency.[34] The results indicate that time aggregation cannot, in the context of the baseline model, account for the empirical findings. Just a little over two percent of plant-year observations in the simulated data involve positive investment and employment contraction, and the average decline in employment among those plants which are contracting is slightly under two percent. These results are illlustrated in the right panel of Figure 3, which plots the model-implied distribution of the log change in employment given positive investment, that is, given a net investment rate in excess of 10 percent. For reference, we show the distribution from the South Korean data next to it.[35]

It is instructive to re-run the annual simulations with a lower threshold for positive investment. Rather than condition on a net investment rate greater than 10 percent, we set the threshold at 0.01. In that case, 18 percent of observations in the simulated data show a year-over-year net decline in employment, given positive investment. However, the average decline in employment among contracting plants is 4 percent, which is less than one-fourth of average in Korean data.

Hence, in the model, contraction does occur, but it is limited. This reason for this is intuitive. Since the model is quarterly, the theory indicates that positive investment – of any magnitude – and negative employment growth do not co-exist at a quarterly frequency. Thus, in the annual panel, we see both occur only if firing later in the year undoes the hiring that accompanied investment earlier in the year. This happens, though, only if the number of hires was small. That in turn means that the increase in investment must have been relatively small, too. So episodes of firing and positive investment are ones where the

[34] At a *quarterly* frequency, positive investment and negative employment growth should not co-exist. In fact, we find that they occur together in only 0.02 percent of sample observations. We view this as a useufl check on the accuracy of the simulations.

[35] The reader will notice that the model-implied distribution is tilted to the left. When (capital) depreciation is set to zero, the distribution becomes more symmetric. Intuitively, since depreciation erases some of the undesired capital but the firm has the same number of workers, a fall in x leaves the firm with relatively more "excess" employment than in the no-depreciation baseline. This makes it more likely that the firm reaches the firing barrier in the model with depreciation.

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. As we discussed, 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.

We now introduce worker attrition. This does not disrupt the theoretical predictions discussed in section 1 – in any given quarter, positive investment and negative employment growth do not occur. But one may suspect that attrition interacts with time aggregation. Specifically, the plant may hire and invest in one quarter, but attrition over the subsequent quarters results in a net decline in employment over the year. To consider this more carefully, we set the attrition rate, δ_n, to be 4 percent per quarter. For the purpose of this exercise, this is likely a fairly generous estimate of attrition. Chang, Nam, and Rhee (2003) actually estimate the separation rate in 1994 in Korea to be no more than 2 percent per quarter.

Results are shown in Table 4. There are two remarks we wish to make in this regard. First, the reader will notice that, since attrition is assumed to be costless, it is now virtually impossible (at least within the baseline model) to generate any inaction in employment. It is possible to remedy this if we assume that there are some jobs which are relatively costless to refill after a separation. But since that merely makes labor even more flexible, it is unclear that such an amendment would affect the bottom line of this analysis.[36]

Second, with regard to the co-movement of capital and labor, the effects of attrition are minimal. The intuition behind this result is straightforward. Attrition does deplete the workforce. But firms also hire more often than otherwise, since attrition makes hiring less costly to reverse. This means that, after a period in which the plant hires and invests, it typically does not undergo quarter after quarter in which attrition simply wears away its workforce – it will tend to hire at some point in the year. This blunts the ability of the model to produce years in which the plant both invests and sees its workforce shrink.[37]

[36] Since we can no longer calibrate the cost of adjusting to the inaction rate, we leave $c_n^- = c_n^+$ at the value used in the prior simulation. Other parameters are updated as needed to reproduce the targeted moments. We have also considered an alternative in which we re-calibrate the cost of adjusting labor so that the probability of not "actively" adjusting equals the unconditional probability of inaction in the model without attrition. A plant is said to not actively adjust its workforce if employment at the end of the year, n, exactly equals the prior year's workforce that survived attrition, $(1-\delta_n)^4 n_{-4}$. This approach requires a much larger cost of adjusting labor in order to to deter hiring even in the face of 4 percent per quarter attrition. But as labor becomes more costly to adjust, we also have to reduce the resale price of capital, or else plants will instead adjust too often along this margin. Thus, in the end, the relative frequency of adjusting labor is retained, and the co-movement of the two factors is very similar to what we show in Table 4.

[37] This is not to say that attrition is unimportant – it accounts for the vast majority of quarterly declines in employment in the simulated data. This role of attrition in the model appears to mirror its role in the Korean labor market in particular. The restrictions on the layoff of permanent workers makes attrition a

3 Extensions to the baseline model

In this section, we discuss a few modifications of the baseline model that may plausibly induce more realistic co-movement of investment and employment growth at the plant level. The theories discussed in this section fall into two broad classes. The first preserves factor-neutral technical change and seeks to account for the data along other lines. The second class introduces factor-biased technology.[38]

3.1 Non-technological explanations

Sticky product prices. In the macroeconomics literature, it has been noted that if the firm's product price is sticky, factor-neutral technical change can be contractionary. If the firm does not wish to lower its price in order to sell additional output, it does not need its current factors given the higher level of TFP. So its factor demand declines. However, in these models, technology is contractionary with respect to both capital and labor (see Basu, Fernald, and Kimball, 2006).[39]

Shifts in factor prices. We have taken the real wage as given and abstracted from factor supply. We have done so because factor price movements in general equilibrium are unlikely to generate qualitatively different joint dynamics at the plant level. If the two physical factors are q-complements in production and if revenue exhibits decreasing returns to k and n jointly, then an increase in the price of either factor (perhaps because of a withdrawal of supply) reduces demand for both factors – factor demands react in the *same* direction. This is straightforward to show in the frictionless model, and it remains true in the model with frictions.[40] Of course, if there is, for instance, an increase in the wage that faces the plant, the optimal capital-labor *ratio* will rise, but capital will not increase absolutely.

Structural change. Recent models of long-run growth have shown that neutral technology improvements in manufacturing can lead to a secular decline in that sector's employment

critical means by which plants shrink their permanent workforces (Grubb, Lee, and Tergeist, 2007).

[38] As we noted in the introduction, our results echo findings of a few earlier studies that analyzed data from the U.S., the Netherlands, and Germany. For this reason, we do not tailor the discussion in this section to the specific institutional features of the economies for which we have data. Our result regarding the co-movement of capital and labor appears to be robust across economies.

[39] This literature has also explored the effects of investment-specific shocks. But in this case, the shocks are typically expansionary despite price stickiness (see Smets and Wouters, 2007).

[40] Appendix C provides a summary of the robustness analysis we have performed that is not reported in detail here. In one exercise, we include exogenous shifts in the (aggregate) price of capital in the model of section 1. As we would anticipate from the frictionless model, these shocks failed to induce positive investment coincidence with employment declines.

share if the price elasticity of demand for the composite manufactured good is sufficiently small. Moreover, the fact that this elasticity of demand at the industry level is so low does not imply a counterfactually high (or any) markup over marginal cost at the individual plants. However, in our data, the vast majority of variation in factor inputs is not due to industry-wide shocks; all of our empirical results hold within narrowly defined sectors. Thus, we interpret the incidence of negative comovement between machinery and employment as largely driven by idiosyncratic shocks to plants. In other words, it does not seem that this pattern is due to a common shock that drives all plants up a steeply sloped industry demand schedule. As a result, the relevant demand elasticity is that which operates at the plant-level. We return to explanations that rely on low plant-level elasticities in section 3.2.

Machinery replacement. When machinery is subject to stochastic breakdown, it is possible to partially de-couple investment from employment adjustment. For instance, suppose that a share of a plant's machinery is "critical" to production in that output falls virtually to zero if this subset of equipment fails (irreparably). Assume these failures are orthogonal to productivity and/or demand (and their histories). This means that there will be quarters in which a firm may invest merely to keep the plant open, even if productivity is so low as to warrant a reduction in employment.

This account of plant-level dynamics has some evidence in its favor. For instance, when we look at transportation equipment (such as trucks used for delivery), we also see that plants sometimes contract employment even as they increase invest in these vehicles. This seems consistent with the notion that the failure of a truck would create such a disruption that replacement is necessary even if productivity is relatively low.

The challenge with respect to this explanation is that it is very hard to assess its quantitative importance. We are not aware of a point estimate (or even a distribution of estimates) as to the share of a plant's machinery that is critical in the sense described above. But this matter deserves further attention.[41]

The lag between order and delivery. The baseline model assumes that new machinery is installed and operated in the period in which it is ordered. But in reality there may be

[41] It should be noted that there is not necessarily a uniform definition of criticality, which makes it hard to apply results from the related maintenance literature. For instance, some research in plant maintenance and reliability treats electricity supply as critical, but we do not, of course, interpret that as a subset of machinery. Furthermore, even if a machine is critical, the plant might have redundancies. This is important if plants are credit-constrained and redundant parts have high prices. In that case, the plant will tend to invest in redundancies when its internal funds are high, which is precisely when it would want to hire. Hence, this model of criticality would not deliver negative comovement between machinery and employment. We wish to thank (without implicating, of course) Andrew Starr (Manufacturing and Materials Department, Cranfield University) for sharing his thoughts on the issue of criticality.

a delay between the order and delivery of new capital goods. If the plant reports investment in the year in which the order is placed, then this delay does not disrupt the implications of the baseline model: orders, and thus recorded investment, should still perfectly predict positive employment growth.

The implications of a delay are more significant if plants record investment in the year in which the machinery is delivered. Suppose machinery is ordered in August 2001 and received in February 2002. If plant productivity deteriorates over the course of 2002, employment may be reduced that year even though positive investment is observed. The delivery lag, in other words, *amplifies time aggregation*. If there were no delay, the delivery would coincide with the order. Thus, productivity would have to decline sufficiently between August 2001 and December 2001 in order to generate a year in which the plant undertakes investment and reduces its labor demand (on net). But if there is a delay, productivity can fall at any point in 2002 and trigger a decline in employment that year even though investment would be recorded.

To investigate this, we generalize the baseline model in a simple way in order to allow for delivery lags. We adopt a Calvo-like mechanism: an order for machinery, regardless of when it was made, is filled with probability λ in any one period. As a result, if an order goes unfilled this period, it takes its place in the prior backlog of unfilled orders, and whole stock is filled with probability λ next period. This assumption greatly simplifies the problem, albeit at the expense of a certain degree of realism.

We now sketch the plant's problem in the presence of a delay. Let κ_{-1} denotes the initial backlog of unfilled orders at the start of the present period. The plant's productivity, x, is revealed, and it chooses its capital and labor demands. With respect to capital, the plant may choose to sell a portion of its installed machinery, k_{-1}. We let $\hat{k}$ denote its machinery stock after any sales have taken place but before orders are placed. If the plant chooses to order machinery, its stock of unfilled orders rises to κ. If these orders arrive, they are delivered in the final instant of the period. At that point, production ensues.

The plant's Bellman equation then becomes

$$V\left(\kappa_{-1}, k_{-1}, n_{-1}, x\right) = \max_{\kappa, \hat{k}, n} \left\{ \begin{array}{c} -wn - \mathcal{C}_n\left(n, n_{-1}\right) - \mathcal{C}_k\left(\hat{k}, \kappa, k_{-1}, \kappa_{-1}\right) \\ +\varrho \mathrm{E}_{x'} \left[\begin{array}{c} \lambda \cdot \left[x^{1-\alpha-\beta}\left(\hat{k}+\kappa\right)^{\alpha} n^{\beta} + V\left(0, \hat{k}+\kappa, n, x'\right)\right] \\ +\left(1-\lambda\right) \cdot \left[x^{1-\alpha-\beta}\hat{k}^{\alpha} n^{\beta} + V\left(\kappa, \hat{k}, n, x'\right)\right] \end{array} \right] \end{array} \right\}.$$

This says that if orders arrive, the plant produces with a capital stock, $\hat{k}+\kappa$. In that case, the backlog of unfilled orders taken into the next period is 0. Otherwise, the plant produces with $\hat{k}$, and takes an unfilled order stock κ into the next period.[42] The cost of adjusting capital, $\mathcal{C}_k\left(\hat{k}, \kappa, k_{-1}, \kappa_{-1}\right)$, is notationally more cumbersome but substantively the same as before. It is given by

$$\mathcal{C}_k\left(\hat{k}, \kappa, k_{-1}, \kappa_{-1}\right) = c_k^+ \left(\kappa - \kappa_{-1}\right)_{1[\kappa > \kappa_{-1}]} + c_k^- \left(k_{-1} - \hat{k}\right)_{1\left[k_{-1} > \hat{k}\right]}.$$

The difference, $\kappa - \kappa_{-1}$, represents the new orders placed this period, and $k_{-1} - \hat{k}$ is the number of units sold.

We solve the model numerically. The calibration of the size of the adjustment costs, the returns to scale, and the stochastic process of productivity (x) does not differ materially from that presented in Table 3. The only free parameter to be pinned down, then, is λ – the probability per quarter that a delivery is made. There is a not great deal of evidence on this, and what we have found pertains to the U.S. rather than Korea. Abel and Blanchard (1988) estimate a delivery lag of 2 to 3 quarters, whereas 2 quarters is on the high end of the range of estimates reported in Carlton (1983). In Table 5, we therefore present results for $\lambda = 1/3$. Note that if $\lambda = 1/3$, then $(2/3)^3 \cong 30$ percent of firms wait longer than 3 quarters for delivery.[43]

The results in Table 5 indicate that delivery lags do not appear to be so large as to account for the majority of the coincidence of positive investment and negative employment growth. As shown in the bottom panel, the share of plant-year observations that involve opposite movements in investment and employment is now 12.5 percent. This is far higher than in the baseline model but still less than one-third of the empirical estimate. The mean decline in log employment, conditional on positive investment, is 3.5 points, which is less than one-fifth what we see in the data. This reflects the difficulty in generating sizable reductions in employment purely from time aggregation. As shown in the middle panel, orders of machinery (at the quarterly frequency) are virtually never accompanied by contractions in labor demand – *desired* employment perfectly co-moves with desired capital. Arguably, a model might have more potential if increases in investment triggered a reduction in desired labor demand. We will take up such a theory in the next subsection.

The disruption of factor adjustment. Cooper and Haltiwanger (1993) provided a model in which machine replacement disrupts production and reduces current labor demand.

[42] Note that we assume the plant pays for the orders when they are placed. This is innocuous in a model with perfect credit markets.

[43] To be clear, when the model is simulated, investment is recorded at the time of delivery. This captures the mis-match between the dates at which employment is adjusted and investment is recorded in the data.

Formally, disruption means that, in the period in which machinery is installed, the marginal product of labor is lower (for any given k). This is intended to represent the idea that production must be stopped or slowed as new machinery is put into place. Further, Cooper and Haltiwanger assume labor is a flexible factor and that machinery is not available for use in production during its installation. It follows that employment will be reduced at the time of installation.

The model of Cooper and Haltiwanger (1993) is deterministic, but it is possible to sustain the argument in a stochastic environment. Suppose the installation of capital disrupts production in the sense that revenue is given by $\left(1-\tau\cdot\mathbf{1}_{[\Delta k>0]}\right)\cdot x^{1-\alpha-\beta}k^{\alpha}n^{\beta}$, where $\tau\in(0,1)$. This says that if investment is undertaken, the marginal product of each factor is degraded by τ for any given (k,n,x).[44] If we retain the assumptions that new capital is not available for use in production and that labor is a frictionless factor, then the choice of labor conditional on $\Delta k>0$ is governed by the static first-order condition, $(1-\tau)\,\beta x^{1-\alpha-\beta}k^{\alpha}n^{\beta-1}=w$, where k is interpreted as the given, start-of-period machinery stock. If the firm was near its Ss band with respect to investment prior to this period, it does not require a large increase in productivity, x, to be induced to invest. If x increases only marginally from last period, the "tax" on the marginal product, τ, may dominate and lead to a fall in labor demand.[45]

In addition to the disruption cost, this set-up differs from the model of section 1 in two ways. The first relates to timing – in section 1, we assume new machinery is available for use in production in the period in which investment is undertaken. This is important because investment is "lumpy" in the presence of a disruption cost.[46] As a result, the large adjustment to k will likely offset the effect of τ on the marginal product of labor, triggering an increase in labor demand. The second difference is that, in the model of section 1, labor is subject to costs of adjusting. In this case, it is much less clear it would be optimal for a firm to reduce employment during the installation period. Even if new machinery is not available this period, the scheduled, discrete increase in capital next period raises the expected marginal value of labor. If layoffs are costly to reverse, the firm has an incentive

[44] A specification like this has been used often in the more recent dynamic factor demand literature. See, among others, Caballero and Engel (1999), Cooper, Haltiwanger, and Willis (2005), and Bloom (2009).

[45] Cooper, Haltiwanger, and Power (1999) incorporate stochastic plant-specific TFP into a machine replacement model. However, their model omits labor adjustment.

[46] The disruption cost, given by $\tau\cdot x^{1-\alpha-\beta}k^{\alpha}n^{\beta}$, is a kind of fixed cost. This is easy to see if new capital does not become operative until next period. In that case, current revenue – and thus, the disruption cost – is independent of the size of the change in the capital stock. But even if capital becomes operative immediately, the disruption cost has a discontuity at the origin, as in $\lim_{\Delta k\to 0}\tau\cdot x^{1-\alpha-\beta}k^{\alpha}n^{\beta}\neq 0$. This means that it rises at an infinitely fast rate when any investment is undertaken. This discontinuity is the unifying feature of fixed costs and gives rise to the economies of scale in adjusting which make lumpy changes optimal.

to retain its workforce.[47]

We now explore this issue numerically. Specifically, we modify the baseline model in two regards. First, a disruption cost associated with capital adjustment is introduced. Second, we assume a one-period delay between the date of installation and the date at which the new capital becomes operative. The question we ask is whether the costly reversibility of labor deters layoffs in the periods in which new capital is installed.

In this modified model, the costs of adjusting must be re-calibrated. We now have to select two parameters with respect to the cost of adjusting capital – the resale price, c_n^-, and the disruption cost, τ. Therefore, we need another moment, in addition to the frequency of investment. We set the resale price of capital to target the frequency of *negative* investment. This leaves τ to target the degree of inaction in investment.[48] We found that it takes only an exceptionally small value of τ to induce a substantial amount of inaction. Yet if τ is set too low, then the model has little chance of engaging the data – the marginal product of labor would not fall by enough to induce declines in employment at the time of installation. To strike a compromise, we fix $\tau = 0.025$, which implies that investment is zero around 70 percent of the time. This inaction rate about 1.5 times that in the data.[49]

Table 5 displays the results. The middle panel of the table shows that the disruption cost does not trigger declines in employment in the quarter in which the installation is undertaken. However, since capital is not operative until the next period, the growth in employment in the installation period is subdued relative to the baseline (see Table 4). This may account for why, at an annual frequency, we do see some instances of employment declines and investment, although the coincidence of the two is still far below that in the data. Suppose the plant fires in quarter one but x then reverses course and by the fourth quarter, it is optimal to hire and invest. Since the new machinery is not operative, the increase in employment is muted in quarter four. This shows through as a larger net reduction in employment for the year.[50]

[47] In a richer model, one might differentiate between a temporary furlough and a permanent severance of the firm-worker match. If the former is allowed, then temporary layoffs would likely coincide with machine installation. It is unlikely that this is the source of the co-movement we see in the data, though. First, since we measure year-over-year changes in our data, this kind of short-term (intra-year) variation in the workforce is unlikely to account for our results. Moreover, as discussed in section 2, there is no evidence that employment declines, conditional on investment, are reversed at a faster rate than otherwise.

[48] We continue to calibrate the cost of labor adjustment to target the frequency of plant-year observations that display zero net employment growth.

[49] Such a low value of τ is in conflict with Cooper and Haltiwanger (Table 5, 2006). But they did not target the adjustment frequency.

[50] On the other hand, if the plant installs machinery in quarter 2, there will be a large increase in employment in the next quarter. This reflects the lumpy change in the marginal product due to (discrete) machinery investment. As a result, Table 5 shows that the average annual change in employment, condi-

3.2 Factor-biased technical change

The coincidence of investment and (net) separations suggests the presence of factor-biased technical change. However, we argue that relatively standard implementations of factor-biased technical change do not induce the incidence of investment and separations that we observe in the data. These production functions do induce, *conditional* on output, a shift away from labor after a technology improvement. However, the scale effect of an increase in technology dominates unless demand is highly inelastic.

Thus, one may view the co-movement of capital and labor in one of two ways. It may be treated as a moment which informs the calibration of the demand schedule and the mark-up, in which case typical estimates of the mark-up used in macroeconomic analysis appear very low. Or, if the implied mark-up clashes with reliable external evidence, it may be interpreted as a challenge to canonical models of technical change.

In this paper, we (at least tentatively) adopt the latter view and so consider an alternative model of factor-biased technology. In this framework, the final good is produced by aggregating over the outputs of many tasks. The firm allocates machinery to a subset of tasks and labor to the remainder. If there an increase in the productivity of machinery in this model, labor is directly replaced in the marginal task by capital. This form of factor bias will enable the model to replicate the coincidence of investment and (net) separations.

Factor-augmenting technology in CES production. We first consider a model of CES production in which there is labor-augmenting technical change. The latter can, in theory, be *labor-saving* and thus capable of inducing a simultaneous increase in capital and reduction in employment. However, for a wide range of parameter values, labor-augmenting innovations are not contractionary with respect to labor.[51]

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,

$$y = \left(\vartheta k^{\frac{\varphi-1}{\varphi}} + (1-\vartheta)\left(\xi n\right)^{\frac{\varphi-1}{\varphi}} \right)^{\frac{\varphi}{\varphi-1}}, \tag{4}$$

where φ is the elasticity of substitution across the factors, ξ is labor-augmenting technical change, and ϑ is related to 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

tional on positive investment, is around 30 log points. We could ratchet down the size of TFP shocks to bring this more closely in line with the data. However, this comes at a price: we would further understate the unconditional dispersion in employment growth.

[51] Capital-augmenting technical change is arguably even less promising as a means of inducing investment and employment contraction, as we discuss below.

faces an isoelastic product demand schedule given by

$$y = \zeta P^{-\epsilon}. \tag{5}$$

where ζ is the demand shifter and P is the price of the plant's product. It follows that the value of the firm becomes

$$\begin{aligned} &\Pi\left(k_{-1}, n_{-1}, \xi, \zeta\right) \\ &= \max\nolimits_{k,n} \left\{ \begin{array}{c} \zeta^{1/\epsilon} y^{\frac{\epsilon-1}{\epsilon}} - wn - \mathcal{C}_k\left(k, k_{-1}\right) - \mathcal{C}_n\left(n, n_{-1}\right) \\ +\varrho \int \int \Pi\left(k, n, \xi', \zeta'\right) f_\xi\left(\xi'|\xi\right) f_\zeta\left(\zeta'|\zeta\right) \mathrm{d}\xi \mathrm{d}\zeta, \end{array} \right\}, \end{aligned} \tag{6}$$

where, for any χ, f_χ is the conditional p.d.f. of χ. 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 the problem is homogeneous of degree one in (k, n, ξ, ζ). This allows us to normalize the factors with respect to ζ, giving $\tilde{n} \equiv n/\zeta$ and $\tilde{k} \equiv k/\zeta$, and recast the problem in terms of $\left(\tilde{k}, \tilde{n}, \xi\right)$. Thus, the policy rules given in section 1 (and plotted in Figure 1) hold here, *conditional on* ξ.

To gauge how well this model matches the plant-level co-movement of capital and labor, it is instructive to consider a simple exercise. Suppose a plant is now investing and firing, so its first-order conditions corresponding to these choices are in effect. Then if ξ is labor-saving, these FOCs must remain in effect after an increase in ξ, implying an increase in k and a decrease in n. To pursue this point, consider the version of (6) with static expectations, i.e., the firm does not anticipate that ζ or ξ will evolve from current levels. In that case, one can differentiate the FOCs for investing and firing to obtain the comparative static, $\mathrm{d} \ln n / \mathrm{d} \ln \xi$. This yields

$$\frac{\mathrm{d} \ln n}{\mathrm{d} \ln \xi} = (\epsilon - 1) - (\epsilon - \varphi) \frac{s}{s+1},$$

where s is the ratio of capital income to labor-related payments (the wage bill plus the amortized cost of firing). This is negative only if

$$\epsilon - 1 < (1 - \varphi) s. \tag{7}$$

This says, in part, that the elasticities of demand and substitution must be sufficiently small. When φ is near zero, the factors are highly complementary, so an increase in the effective supply of labor induces the plant to increase capital relative to labor. This can involve an absolute reduction in labor. However, at the same time, an increase in ξ reduces the marginal cost of production and the product price. This yields an increase in sales. If demand is sufficiently inelastic (ϵ is small), though, this scale effect is muted. As a result, (7) is more likely to hold, and labor declines.

Equation (7) provides a hint as to the difficulty that (4) has in generating contractionary technology improvements. The range of φs that satisfy (7) is increasing in s and declining in ϵ. But even for polar values of s and ϵ, it is difficult to make $\frac{\mathrm{d}\ln n}{\mathrm{d}\ln \xi}$ turn negative. Table 6 illustrates this. To interpret the table, note that (7) implies that the elasticity of substitution, φ, must satisfy $\varphi < \hat{\varphi} \equiv 1 - \frac{\epsilon-1}{s}$. The table shows the values of the threshold, $\hat{\varphi}$, that correspond to values of ϵ listed along rows and s listed along columns. Since φ is bounded below by zero, only the shaded values of $\hat{\varphi}$ are feasible. Hence, if $s = 0.8$, which is the mean among Korean manufacturers, the product elasticity, ϵ, would have to be less 2 to even be consistent with (7). This is well below estimates in the literature.[52]

Of course, the ratio s is likely to vary across plants. Perhaps the coincidence of positive investment and (net) separations in the data reflects the presents of plants that are relatively capital intensive. Unfortunately, it is hard to investigate this in our data since we do not have detailed information on capital income.[53]

What we can say is that, if s were higher, one must still invoke what seem like low values of (ϵ, φ) to make (7) hold. For example, suppose $s = 1.6$, or twice its empirical average. We are not aware of plant-level estimates of φ, but a value of 0.1 brings us close to Leontief and arguably serves as a reasonable lower bound. In that case, one needs values of ϵ below 2.44 in order to make (7) hold. This is less than the estimates in Broda and Weinstein (2006, Table IV, "Median" row, "TSUSA/HTS" column). Moreover, in this model, an ϵ of 2.44 implies a markup of almost 70 percent. This conflicts with recent evidence from Feenstra and Weinstein (2010), who estimate a richer translog demand system and find an implied median mark-up of around 30 percent.[54]

We now demonstrate that our intuition gleaned from (7) is highly indicative of the behavior of the dynamic factor demand model. To do this, we set $\epsilon = 2.44$, $\varphi = 0.1$, and $s = 1.6$ and simulate plant-level data on investment and employment growth.[55] Table 7 re-

[52]According to the model, we would like to measure the denominator in s by summing the wage bill and the amortized cost of firing. However, the data do not include the latter. Thus, the empirical estimate of $s = 0.8$ likely *overstates* the value of s that is relevant for evaluating (7).

[53]For the U.S., the Bureau of Labor Statistics does provide industry-level data on rental prices. These data reveal that the average value of s is *less than one* in 14 of 16 three-digit NAICS manufacturing industries. This limits the degree to which plausible variation in s could account for the degree of negative co-movement between capital and labor documented by Sakelleris (2000) on U.S. data.

[54]Both of these papers use detailed data on U.S. imports. In particular, the Broda and Weinstein estimate is based on 10-digit Harmonized Tariff System data. We highlight these estimates, since the more disaggregated data are arguably more indicative of the elasticities faced by individual plants.

[55]For this exercise, we omit capital disruption costs. Thus, we re-calibrate the resale price of capital and

ports the results. Positive investment and negative employment growth hardly ever coexist in the simulated data.

Before we leave this subsection, we note that a model of capital-augmenting technical change is even less effective at engaging the data. To see this, suppose production is now given by

$$y = \left(\vartheta\left(\xi k\right)^{\frac{\varphi-1}{\varphi}} + \left(1-\vartheta\right) n^{\frac{\varphi-1}{\varphi}}\right)^{\frac{\varphi}{\varphi-1}}. \tag{8}$$

Rearranging this, we may write $y = n\left(\vartheta\Gamma^{\frac{\varphi-1}{\varphi}} + \left(1-\vartheta\right)\right)^{\frac{\varphi}{\varphi-1}}$, where $\Gamma \equiv \frac{\xi k}{n}$ is the effective capital-labor ratio. This indicates that the effect of an increase in ξ on labor demand will hinge on the reaction of the marginal revenue product of labor, $MRPL$, to a change in Γ. Specifically, the firm's first-order conditions imply[56]

$$\frac{\mathrm{d}\ln n}{\mathrm{d}\ln\xi} = \left(\epsilon-\varphi\right)\frac{s}{s+1} \propto \frac{\mathrm{d}\ln MRPL}{\mathrm{d}\ln\Gamma}. \tag{9}$$

If $\epsilon > \varphi$, $\frac{\mathrm{d}\ln n}{\mathrm{d}\ln\xi}$ is unambiguously positive.

The condition on $\epsilon - \varphi$ is intuitive, since $\epsilon > \varphi$ implies the factors are q-complements. This accounts for why, under $\epsilon > \varphi$, an increase in the effective supply of capital raises labor demand. Unfortunately, it is hard to assess the restriction $\epsilon > \varphi$, since estimates of the elasticity of substitution φ, conditional on the form of (8), are not available at the plant level. But Raval (2012) does provide estimates off U.S. county-level data for each two-digit manufacturing sector. Virtually all of his estimates are less than 1. Therefore, acknowledging the obvious caveats, we view the restriction $\epsilon - \varphi > 0$ as a reasonably safe one to make.

Investment-specific technical change & skill bias. Thus far, we have assumed only two factors. Within the skill-biased technical change, literature, though, it is common to allow for skilled and unskilled labor, such that the capital and skilled labor are gross complements whereas capital and unskilled labor are gross substitutes. Krusell, Ohanian, Rios-Rull, and Violante (2000) consider this kind of production function. Technical change in their model takes the form of advances in the rate at which investment is transformed into installed capital.

There are two reasons we do not believe this model is consistent with our data. First, the model would lead us to expect that employment declines are concentrated among low-skilled workers. But the evidence for this is far from clear. In the Korean census, plants record the number of production and non-production workers. In the literature on skill-biased

the cost of labor adjustment to target the individual frequences of adjusting.

[56] We again derive the comparative static under the assumption of static expectations.

technical change, non-production status is often treated as a proxy for skill.[57] Therefore, we re-compute the share of plant-year observations that involve both positive investment and declines in non-production workers. This is 35.5 percent, which is only slightly less than the baseline estimate reported in Table 1. Moreover, conditional on positive investment and negative non-production employment growth, the average contraction in non-production workers is not any smaller than the decline among the general workforce reported in Table 1.[58]

These results are not that surprising in light of what we have learned about the impact of technology on the workplace. The tasks performed by clerical and administrative workers, who would be classified as non-production employees in our data, have become increasingly codified and performed by machines (Acemoglu and Autor, 2010). This suggests that it is not only workers on the factory floor for whom machines may substitute.[59]

The second reason we doubt this model is consistent with the plant-level is that, even if it induces declines in unskilled labor, it is very unlikely to generate *plant-wide* declines in labor demand given plausible parameter values. This is straightforward to show in the model of Krusell, et al, where factors are assumed to be perfectly flexible.

Suppose there are skilled, s, and unskilled, u, workers. Production is given by

$$y = \left(\vartheta_u u^{\chi} + (1-\vartheta_u)\,\Xi(k,s)^{\chi}\right)^{1/\chi}, \tag{10}$$

where $\Xi(k,s) \equiv \left(\vartheta_k k^{\rho} + (1-\vartheta_k)\, s^{\rho}\right)^{1/\rho}$ is the CES bundle of capital and skilled labor and $\vartheta_u, \vartheta_k \in (0,1)$. Based on the estimates in Krusell, et al, we assume that $\chi < 1$, which implies an elasticity of substitution, $\varphi_{\chi} \equiv \frac{1}{1-\chi}$, between unskilled labor and the capital-skill bundle that is greater than one. We also assume $\rho < 0$, which implies an elasticity of substitution, $\varphi_{\rho} \equiv \frac{1}{1-\rho}$, between capital and skill that is less than one. If demand is given by equation (5), then revenue is $y^{\frac{\epsilon-1}{\epsilon}}$. Hence, the firm maximizes

$$\sum_{s=0}^{\infty} \beta^s \left\{ \begin{array}{c} \left[\vartheta_u u_{t+s}^{\chi} + (1-\vartheta_u)\left(\vartheta_k k_{t+s}^{\rho} + (1-\vartheta_k)\, s_{t+s}^{\rho}\right)^{\frac{\chi}{\rho}}\right]^{\frac{1}{\chi}\frac{\epsilon-1}{\epsilon}} \\ -w_u u_{t+s} - w_s s_{t+s} - i_{t+s} \end{array} \right\},$$

[57] See, e.g., Berman, Bound, and Machin (1998) and Berman and Machin (2000).

[58] Nonetheless, there has been an upward trend in the non-production worker share in our data. Fuentes and Gilchrist (2005) also find such an upward trend in Chile even though the non-production worker share in their dataset does not tend to increase in periods of investment spikes. They argue that the source of the secular trend was a shift in favor of R&D-related personnel after trade liberalization.

[59] The evidence for the prevalence of computing equipment is typically taken from the U.S. and European countries. But South Korea is a relatively advanced economy by this measure. If we take the volume of exports and imports of computer equipment (per capita) as indicative of pervasiveness of this technology, then Korea ranks on par with several Western European and North Atlantic economies (see OECD, 2008).

subject to $k_{t+s} = (1-\delta)\,k_{t+s-1} + q_{t+s} i_{t+s}$. Here, q is the investment-specific technology shock.

To get some insight into the model, fix $q_t = q_{t+j}$ for $j \geq 1$ and consider a once-and-for-all increase in q in time t. We first compute the change in the demand for unskilled labor in response to this shift in q. Appendix B shows that $\frac{\mathrm{d}\ln u_t}{\mathrm{d}\ln q_t} < 0$ if and only if the elasticity of substitution between capital (and skill) and unskilled labor exceeds the elasticity of product demand, $\varphi_\chi > \epsilon$. In other words, the scale effect – the increase in the demand for labor derived from the increase in output – must be dominated by the incentive to substitute capital for unskilled labor. In what follows we assume this holds.

We then compute the response of plant-wide employment, $n \equiv s + u$, to this shift in q. Investment-specific technology is labor-saving if $\frac{\mathrm{d}\ln n_t}{\mathrm{d}\ln q_t} < 0$. Appendix B shows that this holds only if

$$\mu < \frac{\frac{1}{2}\left(\varphi_\chi + \varphi_\rho\right) - \epsilon}{\varphi_\chi - \epsilon},$$

where $\mu \equiv \frac{\partial \ln y}{\partial \ln u} \in (0,1)$ is the elasticity of output with respect to unskilled labor. By assumption, $\varphi_\chi > \epsilon$. However, just to make sure the numerator on the right side is positive (so that there are at least *some* values of μ consistent with $\frac{\mathrm{d}\ln n_t}{\mathrm{d}\ln q_t} < 0$), a stronger condition is needed. Since $\chi \in (0,1)$ and $\rho < 0$, then $\varphi_\chi > \varphi_\rho$, in which case we require $\varphi_\chi > \frac{1}{2}\left(\varphi_\chi + \varphi_\rho\right) > \epsilon$. To assess this, suppose $\chi = 2/3$ and $\rho = -0.5$, as Krusell, et al found. Then $\frac{1}{2}\left(\varphi_\chi + \varphi_\rho\right) = \frac{11}{6} < 2$: this is the *upper bound* on ϵ. As noted above, elasticities of demand below 2 fall well outside the range in the literature.

Machinery replaces labor in tasks. There has been growing interest in models in which machinery directly replaces labor in tasks. As we will show, this model has a greater capability of generating the incidence of investment and separations observed in the data. Our analysis here is based on the framework of Acemoglu (2010) (who in turn adapts the model of Zeira (1998)).

The plant produces a single good that is subject to the demand schedule (5). The good is produced by aggregating the outputs of a continuum of tasks. A task is indexed by v on the unit interval. We assume plant output is then given by the CES aggregate,

$$y = \left[\int_0^1 y(v)^{\frac{\varphi-1}{\varphi}}\,\mathrm{d}v\right]^{\frac{\varphi}{\varphi-1}}, \tag{11}$$

where $y(v)$ is the output of task v. A task may be performed by labor or machinery.[60]

[60] That is, we imagine a single type of worker and a single machine. This same machine is used in each

Specifically, one unit of output in task v can be produced with $\lambda(v)$ units of labor or $\varkappa(v)$ units of machinery:

$$y(v) = \begin{cases} \frac{n(v)}{\lambda(v)}, & \text{or} \\ \frac{k(v)}{\varkappa(v)}. \end{cases} \tag{12}$$

We assume tasks may be ordered from 0 to 1 such that tasks at which machines are relatively productive are ordered first. Therefore, the relative productivity of labor, $\theta(v) \equiv \frac{1/\lambda(v)}{1/\varkappa(v)} = \frac{\varkappa(v)}{\lambda(v)}$, is increasing in v.

The plant's optimization problem can be solved in two steps. First, given total workers n and machinery k, the firm chooses which tasks to perform with capital and which to perform with labor. It also decides the quantities of each factor to allocate to its respective tasks. These are static problems. Second, the firm solves a dynamic problem to select n and k, subject to the adjustment costs described in equation (1). We now briefly summarize the first step.

Formally, the allocation of factors across tasks solves the following problem. (The details of the solution are provided in Appendix B.) The monotonicity of $\theta(v)$ suggests the existence of a single crossing, v^*, such that all tasks $v \leq v^*$ are automated and all other tasks are performed by labor. Taking this as given for the time being, the firm's problem is to decide the quantities of each factor to allocate per task in order to maximize revenue. Hence, given (5) and (11), the firm solves

$$\max_{\{y(v)\}_{v=0}^{v=1}} \zeta^{1/\epsilon} \left[\int_0^1 y(v)^{\frac{\varphi-1}{\varphi}} \mathrm{d}v \right]^{\frac{\epsilon-1}{\epsilon}\frac{\varphi}{\varphi-1}}$$

subject to the resource constraints,

$$k = \int_0^1 k(v)\,\mathrm{d}v = \int_0^{v^*} \varkappa(v)\, y(v)\,\mathrm{d}v$$

$$n = \int_0^1 n(v)\,\mathrm{d}v = \int_{v^*}^1 \lambda(v)\, y(v)\,\mathrm{d}v.$$

Here, $n(v)$ denotes the quantity of labor allocated to task v, and $k(v)$ represents the number of machines allocated to task v.

The solution to this problem gives the optimal quantity $y(v)$ of each intermediate in terms of k , n, and v^*. We are then able to rewrite the plant-wide production function (11) exclusively in terms of the total stocks, k and n, and the threshold, v^*. Specifically, one can

---

task, but, as we will see, it is relatively more productive in some tasks.

show that

$$y = \left[A(v^*)^{\frac{1}{\varphi}} k^{\frac{\varphi-1}{\varphi}} + B(v^*)^{\frac{1}{\varphi}} n^{\frac{\varphi-1}{\varphi}}\right]^{\frac{\varphi}{\varphi-1}}, \tag{13}$$

where the terms, $A(v^*)$, and $B(v^*)$, are defined by

$$\begin{array}{l} A(v^*) \equiv \int_0^{v^*} \left(\frac{1}{\varkappa(v)}\right)^{\varphi-1} \mathrm{d}v \\ B(v^*) \equiv \int_{v^*}^1 \left(\frac{1}{\lambda(v)}\right)^{\varphi-1} \mathrm{d}v \end{array}.$$

To understand these terms, recall that $\frac{1}{\varkappa(v)}$ is output per machine in task v. Hence, $A(v^*)$ is an aggregate over the average products of machinery on the tasks to which machinery is applied. In this sense, it may be interpreted as an index of capital productivity. The same idea applies, of course, to $B(v^*)$.

The indexes, A and B, differentiate (13) from the standard models of factor-augmenting technical change. For instance, if output is produced according to (8), then an increase in capital-augmenting technology does not imply a degradation, in *absolute terms*, in the contribution of labor to output. In other words, there is no decline in the analogue of B in that model. However, when production is given by (13), an increase in the threshold, v^*, leads to a replacement of labor in tasks, which implies a greater contribution per unit of (plant-wide) capital, k, *and* a smaller contribution per worker. This deterioration in labor productivity amplifies the effective bias of technical change.

Thus far, we have taken v^* as given, but it is straightforward to determine the threshold. Cost minimization implies that the firm uses machinery for task v if the marginal cost of machinery is less. The marginal cost is evaluated using the shadow price of capital to the firm, conditional on the total stock k. Since this shadow price is declining in k, it is not surprising that the threshold task is increasing in the capital-labor ratio. Formally, as demonstrated in Appendix B, v^* satisfies

$$\frac{k}{n} = \theta(v^*)^{\varphi} \frac{A(v^*)}{B(v^*)} \equiv \Theta(v^*) \tag{14}$$

The right side of (14) is increasing in v^*, implying a single-crossing. Equation (14) is intuitive, as it says that the optimal allocation of factors across tasks requires that the firm executes a greater share of tasks with the more abundant factor.

It should be noted at this point that the solution (14) assumes a frictionless reallocation of factors within plant across tasks. This abstracts from a number of issues, but we do not believe this omission invalidates the application of (13) to problems of plant-level dynamics. For instance, machinery is likely specific, to some degree, to the task. This raises the question of whether we ought to require the plant to purchase new machinery when capital is deployed to new tasks. This is arguably realistic, but it is not clear to us that it would significantly affect the model's principal predictions. In addition, the model allows the manager to reallocate workers from tasks below v^* to tasks above the threshold. Of course, in reality, a worker on one task may not have the skills to perform others. But this would likely amplify the depressive effects of capital-specific technology on labor demand.[61]

To study the short-run dynamics of factor demand, we now introduce exogenous variation in the (relative) productivity of machines. We do this in a way that preserves continuity with the standard model of capital-augmenting technology, (8). This allows us to more easily compare results across the two specifications. Assume $\varkappa(v)$ takes the form

$$\varkappa(v) = \xi^{-1}\bar{\varkappa}(v), \tag{15}$$

where $\bar{\varkappa}$ is a deterministic, increasing function of v, and ξ is an innovation to machinery productivity. Note that an increase in ξ raises the productivity of machinery on *all* tasks.

Equation (15) implies that $A(v^*)$ in becomes $\xi^{\varphi-1}\bar{A}(v^*)$, where $\bar{A}(v^*) \equiv \int_0^{v^*} \bar{\varkappa}(v)^{1-\varphi}$. Likewise, we may write $\Theta(v^*)$ in (14) as $\xi^{-1}\bar{\Theta}(v^*)$, where $\bar{\Theta}(v^*) \equiv \left(\frac{\bar{\varkappa}(v^*)}{\lambda(v^*)}\right)^{\varphi} \frac{\bar{A}(v^*)}{B(v^*)}$. It then follows that

$$v^* = \bar{\Theta}^{-1}\left(\frac{\xi k}{n}\right). \tag{16}$$

We see that a rise in ξ motivates the substitution of machines for labor in a greater share of tasks.[62]

Equations (13) and (16) imply that the production function for plant-wide output is now

$$y = \left[\mathfrak{a}\left(\frac{\xi k}{n}\right)^{\frac{1}{\varphi}} (\xi k)^{\frac{\varphi-1}{\varphi}} + \mathfrak{b}\left(\frac{\xi k}{n}\right)^{\frac{1}{\varphi}} n^{\frac{\varphi-1}{\varphi}}\right]^{\frac{\varphi}{\varphi-1}}. \tag{17}$$

where $\mathfrak{a}\left(\frac{\xi k}{n}\right) \equiv \bar{A}\left(\bar{\Theta}^{-1}\left(\frac{\xi k}{n}\right)\right)$ and $\mathfrak{b}\left(\frac{\xi k}{n}\right) \equiv B\left(\bar{\Theta}^{-1}\left(\frac{\xi k}{n}\right)\right)$ are increasing and decreasing functions, respectively. With (17) in hand, it is straightforward, in principle, to solve the dynamic factor demand problem. We merely replace the standard CES production function in (6) with (17).

To understand the implications of (17), it is instructive to first analyze the static version

[61] For a detailed analysis of a task-based model in which workers have heterogeneous skills, see Acemoglu and Autor (2010).

[62] This discussion omits shifts in labor-augmenting productivity. But if both factor-augmenting shocks take the form in (15), then one may normalize with respect to the labor productivity shock. To be precise, suppose $\lambda(v) = \Xi^{-1}\bar{\lambda}(v)$, where Ξ is labor-augmenting productivity. Then B becomes $B(v^*) \equiv \Xi^{\varphi-1}\bar{B}(v^*)$ and (14) implies that the threshold task, v^*, is given implicitly by $\frac{(\xi/\Xi)\cdot k}{n} = \bar{\Theta}(v^*)$. Hence, it is only the ratio of ξ to Ξ that affects the selection of the threshold task.

of the model, though. As Acemoglu (2010) stressed, improvements in capital-augmenting technology are contractionary with respect to employment only if the addition of machinery reduces the marginal revenue product of labor. We noted that this is very unlikely in the standard model, given by (8). But in a model of tasks, it is feasible, because an increase in the productivity of machinery – that is, an increase in the effective supply of machinery, ξk – reduces, in absolute terms, the contribution of labor to output and so lessens the demand for labor.

It is particularly easy to see this point if we consider a special case in which $\mathfrak{a}$ is constant. This helps to isolate the effect on labor productivity of shifts in the threshold task. Let $\Gamma \equiv \frac{\xi k}{n}$ again denote the ratio of effective capital to labor. Appendix B shows that the effect of a log-point increase in Γ on the (log of the) marginal revenue product of labor is proportional to

$$\left(1 - \Upsilon\left(\varepsilon_{\mathfrak{b}\Gamma}\right)\right)\varepsilon_{\mathfrak{b}\Gamma} + (\epsilon - \varphi)\frac{s}{1+s}, \tag{18}$$

where, as we will discuss shortly, $\Upsilon\left(\varepsilon_{\mathfrak{b}\Gamma}\right) \leq 0$. Notice that if $\mathfrak{b}$ is independent of Γ, then this collapses to the comparative static () from the standard model; in that case, technology improvements expand labor demand as long as $\epsilon > \varphi$. Otherwise, if $\mathfrak{b}'(\Gamma) < 0$, then an increase in the productivity of machinery can depress labor demand by triggering the replacement of labor with machinery in the marginal task. As for the term $\Upsilon\left(\varepsilon_{\mathfrak{b}\Gamma}\right)$, this is related to the curvature of $\ln \mathfrak{b}$ – in particular, it reflects the elasticity of $\varepsilon_{\mathfrak{b}\Gamma}$ with respect to Γ. Thus, if $\mathfrak{b}$ is a log-linear function of Γ, then Υ is zero. Otherwise, if $\ln \varepsilon_{\mathfrak{b}\Gamma}$ declines in $\ln \Gamma$, then one may show that $1 - \Upsilon\left(\varepsilon_{\mathfrak{b}\Gamma}\right) > 1$, and the bias against labor is amplified further. Intuitively, if $\varepsilon_{\mathfrak{b}\Gamma}$ is not only negative but grows larger in absolute value as Γ increases, the incentive to switch away from labor is greater. This is the case we will consider in the quantitative exercise below.

For the present study, we only wish to illustrate the quantitative potential of (17). Thus, our strategy is to specify flexible parametric forms for $\mathfrak{a}$ and $\mathfrak{b}$ and to select a particular calibration that engages the data reasonably well. We hope in future research to undertake a more detailed estimation of the model. To that end, the flexibility of the functional forms of $\mathfrak{a}$ and $\mathfrak{b}$ will be valuable, as they will allow the data to select the best parameterization.

We use an affine transformation of the beta function for both $\mathfrak{a}$ and $\mathfrak{b}$. Specifically, for $\mathfrak{a}$, we assume

$$\mathfrak{a}(\Gamma) = \kappa_0 + \bar{\kappa}\mathcal{B}_{\frac{\Gamma}{\Gamma+1}}\left(\kappa_1, \kappa_2\right), \tag{19}$$

where $\kappa_0, \kappa_1, \kappa_2, \bar{\kappa} > 0$ and $\mathcal{B}$ is the (incomplete) beta function.[63] The top panel of Figure

[63] The incomplete beta function is defined as $\mathcal{B}_x\left(\kappa_1, \kappa_2\right) \equiv \int_0^x t^{\kappa_1 - 1}(1-t)^{\kappa_1 - 1}\,\mathrm{d}t$, with $x \in [0,1]$.

4 illustrates a few examples of $\mathfrak{a}(\Gamma)$. Note that this may be globally concave or convex. It also may have inflection points. The flexibility of the beta function recommends it for $\mathfrak{b}(\Gamma)$, as well. In particular, we assume

$$\mathfrak{b}(\Gamma) = \eta_0 - \bar{\eta}\mathcal{B}_{\frac{\Gamma}{\Gamma+1}}(\eta_1, \eta_2). \tag{20}$$

Examples of $\mathfrak{b}$ are given in the bottom panel of the figure.[64]

A few considerations guide the calibration of (19) and (20). First, if $\mathfrak{a}(\Gamma)$ increases too sharply (in the relevant region of Γ), then the demand for capital can be upward-sloping. This is not too surprising, in hindsight. As Γ rises, capital does run into diminishing returns for a *given* v^*. But at the same time, the additional capital is deployed to more tasks, which raises the contribution of capital to output. If $\mathfrak{a}(\Gamma)$ is very steep, this latter effect is large, and the increase in the productivity of capital triggers an increase in capital demand. We wish to retain downward-sloping factor demand schedules, if only to preserve continuity with the standard model. This guided the choice of the pair (κ_1, κ_2). There were, in the end, several possibilities, and we have found that, within the range we considered, the precise choice did not have much effect on the plant-level comovement of capital and labor. We settled on $(\kappa_1, \kappa_2) = (1, 3)$. For the sake of symmetry – we did not have much else to guide us – we then set $(\eta_1, \eta_2) = (1, 3)$, too.

Second, the intercepts κ_0 and η_0 clearly affect the capital share of revenue. However, since we work in partial equilibrium, the wage can be calibrated to target the desired capital share. Thus, each of the intercepts can be normalized to 1.

Third, for given shape parameters (κ_1, κ_2) and (η_1, η_2), the slopes $\bar{\kappa}$ and $\bar{\eta}$ govern the rate at which labor is overtaken in tasks as Γ climbs higher. We normalize $\bar{\kappa} = 1$ and then identify a value of $\bar{\eta}$ that induces a realistic probability of employment decline conditional on positive investment. This implied $\bar{\eta} = 1.433$.

Lastly, we must calibrate the stochastic process of the factor-biased technology shock, ξ. In order to retain comparability with the treatment of technical change in the standard models of section 3.1, we assume that ξ follows a simple geometric AR(1) process. However, we have little to guide the calibration of the persistence of the process and the size of the innovations. We select illustrative values. The AR(1) coefficient is set to 0.9, and the standard deviation of the innovation is 0.075.[65]

[64] Clearly, there is a tension between two objectives: the desire for flexibility in $\mathcal{A}$ and $\mathcal{B}$ and the ability to identify these functions. As written, with seven parameters (one of the intercepts may be normalized), these functions are under-identified. At this stage, then, we consider only an illustrative calibration. We are at work on a variant of (17) that is more parsimonious.

[65] It seem awkward to think of negative shocks to machinery productivity in this context, since the trend,

Table 7 reports the results. They appear to largely coincide with the intuition gleaned from (18). The share of observations that involve a decline in employment, conditional on positive investment, now nearly replicates that observed in the data. The model also reproduces the size of the decline in these episodes. At the same time, the presence of a factor-neutral (demand) shock, ζ, allows the model to induce, at other instances, realistically large expansions in employment conditional on positive investment. Figure 5 illustrates the results graphically. It plots the distribution of employment growth conditional on positive investment in both the data and model. This represents a stark improvement relative to the baseline model (see Figure 3).

4 Macroeconomic implications

The plant-level data have pushed us in the direction of a model in which machinery directly replaces labor in production tasks. In this section, we briefly demonstrate a macroeconomic implication of this result. Specifically, we amend the model based around (17) to allow for a stochastic investment subsidy. A reduction in the after-tax price of machinery will stimulate investment. Hence, the reaction of labor demand to the subsidy rests on how the marginal revenue product of labor responds to a shift in the desired capital-labor ratio. But this is governed by (18). Since $\mathfrak{b}' < 0$, the increase in machinery will displace some labor from tasks and potentially depress aggregate labor demand.

A stochastic investment tax policy is incorporated in a simple way. Based on evidence from Cummins, Hasset, and Hubbard (1994), we calibrate a process for the subsidy such that it fluctuates between 0 and 10 percent and changes, on average, once every 2.5 years. Once we have solved for the capital and labor demand functions, we then calculate the impulse response of aggregate employment to a five percentage-point increase in the subsidy. Specifically, we assume the subsidy rises from 2.5 to 7.5 percent and remains high for 10 quarters, before it returns to its original level.

The impulse response of employment is plotted in Figure 6. As we anticipated, aggregate employment declines after the increase in the subsidy. The presence of capital adjustment costs appears to elongate the response, so the decline continues well after the initial impulse. After 10 quarters, the subsidy is reduced, and employment begins to climb.

This is of course only illustrative, in no small part because it is carried out in partial

in reality, has been toward the steady replacement of labor by machinery. But remember (footnote 61) that ξ is more accurately interpreted as the productivity of machinery relative to labor. We assume that there may be innovations to how labor is used in production that imply at least short run declines in ξ.

equilibrium. Nonetheless, it demonstrates the potential of the plant-level data, and the model suggested by those data, to inform our analysis of macroeconomic policy shocks.

5 References

Abel, Andrew and Olivier Blanchard (1988). "Investment and Sales: Some Empirical Evidence," in William A. Barnett, Ernst R. Berndt, and Halbert White (eds.), *Dynamic Econometric Modelling*, Cambridge University Press, pp. 269-296.

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

Acemoglu, Daron and David Autor (2011). "Skills, Tasks and Technologies: Implications for Employment and Earnings," *Handbook of Labor Economics*, 4(B), pp. 1043-1171.

Acemoglu, Daron (2010). "When Does Labor Scarcity Encourage Innovation?," *Journal of Political Economy*, Vol. 118, No. 6, pp. 1037-1078.

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.

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.

Berman, Eli, John Bound and Steve Machin (1998). "Implications of Skill-Biased Technological Change: International Evidence," *Quarterly Journal of Economics*, Vol. 113, pp. 1245-79.

Berman, Eli and Steve Machin (2000). "Skill-Biased Technology Transfer Around the World," *Oxford Review of Economic Policy*, Vol. 16, No. 3, pp. 12-22.

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.

Broda, Christian and David Weinstein (2006). "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, Vol. 121, Issue 2 (May), pp. 541-585.

Carlton, Dennis (1983). "Equilibrium Fluctuations When Price and Delivery Lags Clear the Market," *Bell Journal of Economics*, (Autumn).

Chang, Yongsung, Jaeryang Nam, and Changyong Rhee (2004). "Trends in Unemployment Rates in Korea: A Search-Matching Model Interpretation," *Journal of the Japanese* *and International Economies*, Vol. 18, No. 2, pp. 241-263.

Corbo, Vittorio and Andres Solimano (1991). "Chile's Experience with Stabilization, Revisited," World Bank WPS 579.

Cooper, Russell and John Haltiwanger (1993). "The Macroeconomic Implications of Machine Replacement: Theory and Evidence," *American Economic Review*, Vol. 83, pp. 360-382.

Cooper, Russell, John Haltiwanger, and Laura Power (1999). "Machine Replacement and the Business Cycle: Lumps and Bumps," *American Economic Review*, Vol. 89, pp. 921-46.

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 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.

Edge, Rochelle (2000). "The Effect of Monetary Policy on Residential and Structures Investment Under Differential Project Planning and Completion Times," International Finance Discussion Papers 671.

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).

Feenstra, Robert and David Weinstein (2010). "Globalization, Markups, and the US Price Level," NBER WP #15749.

Fuentes, Olga and Simon Gilchrist (2005). "Skill-biased technology adoption: Evidence for the Chilean manufacturing sector," mimeo.

Grubb, David, Jae-Kap Lee, and Peter Tergeist (2007). "Addressing Labour Market Duality in Korea," OECD Social, Employment, and Migration Working Papers No. 61.

Khan, Aubhik and Julia Thomas (2011). "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity," mimeo.

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.

Krusell, Per, Lee Ohanian, José-Víctor Ríos-Rull and Giovanni Violante (2000). "Capital-skill complementarity and inequality," *Econometrica*, Vol. 68, No. 5.

Letterie, Wilko, Gerard Pfann, and J. Michael Polder (2004). "Factor adjustment spikes and interrelation: an empirical investigation," *Economics Letters* Vol. 85, pp. 145–150.

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.

OECD (2008). *Information Technology Outlook*.

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

Polder, Michael and Sher Verick (2004). "Dynamics of Labour and Capital Adjustment - A Comparison of Germany and the Netherlands," IZA DP No. 1212.

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

Raval, Devesh (2012). "Non Neutral Technology and the Microeconomic Production Function," mimeo.

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.

Veracierto, Marcelo (2002). "Plant Level Irreversible Investment and Equilibrium Business Cycles", *American Economic Review*, Vol. 92 (March), pp. 181-97.

Veracierto, Marcelo (2008). "Firing Costs and Business Cycle Fluctuations", *International Economic Review*, Vol. 49, No. 1 (February).

TO BE ADDED:

Appendix A: Results from Chilean data

Appendix B.1: Baseline simulation model

Appendix B.2: Details of derivations

Appendix C: Robustness analysis

Figure 1: The optimal policy

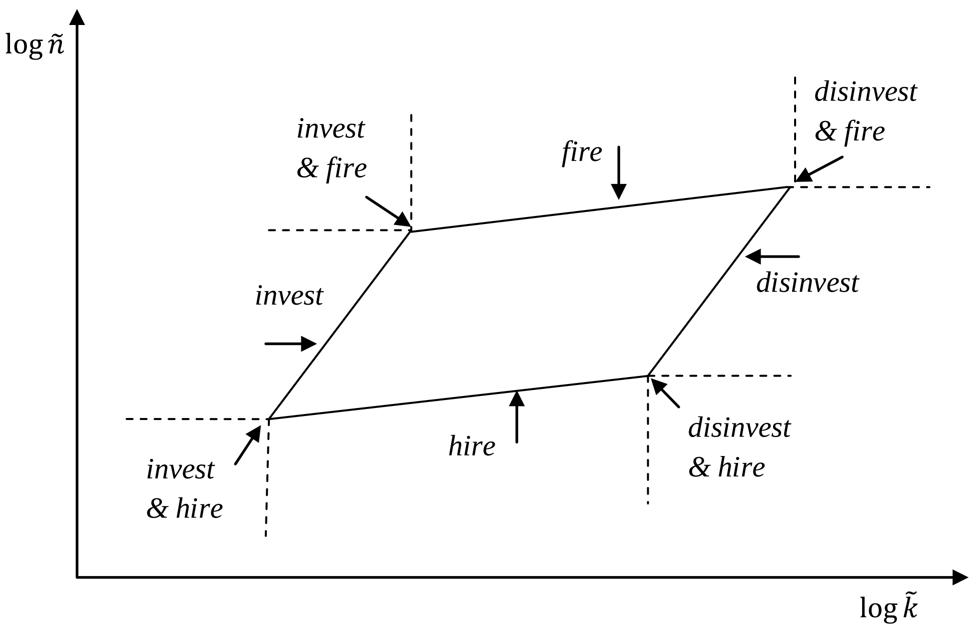

Figure 2: The joint dynamics of capital and employment

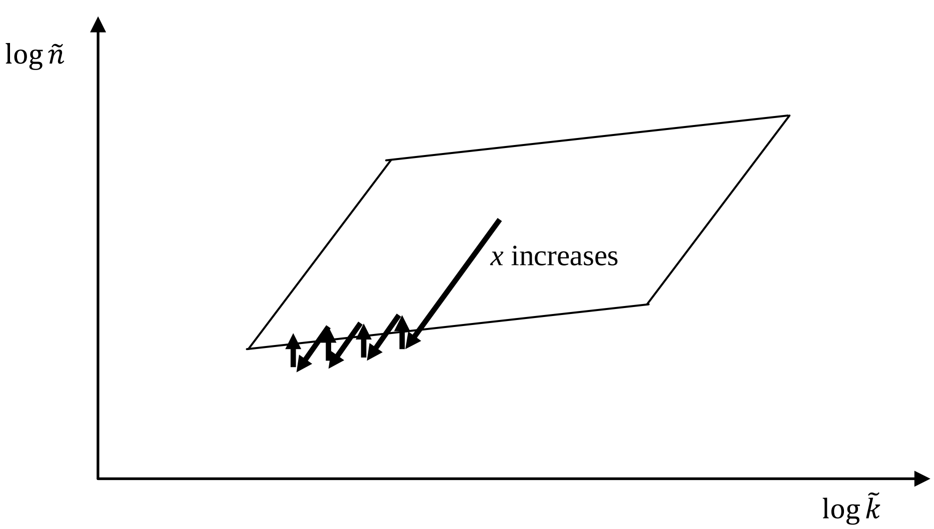

Figure 3: The distribution of employment growth in the data and baseline model

A. Data *B. Baseline model*

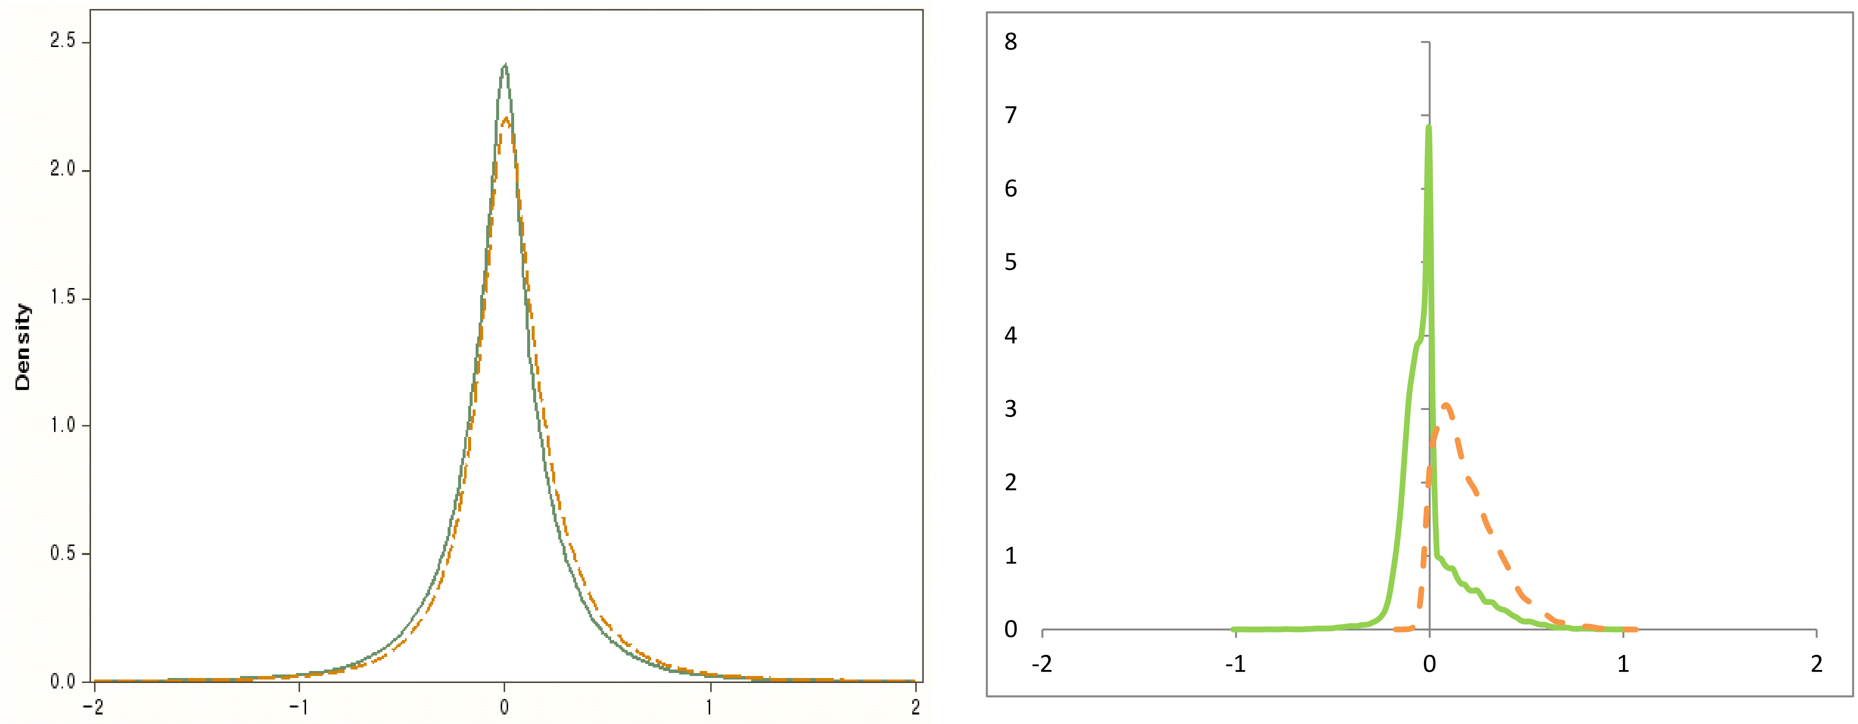

NOTE: The left panel shows the distribution of employment growth in Korean data. The solid (green) line is the kernsel-density estimate of the unconditional distribution, and the dashed (orange) line is the distribution conditional on an investment rate in excess of 10 percent. The right panel shows the distributions in the baseline model of section 1.

Figure 4: Examples of the beta function

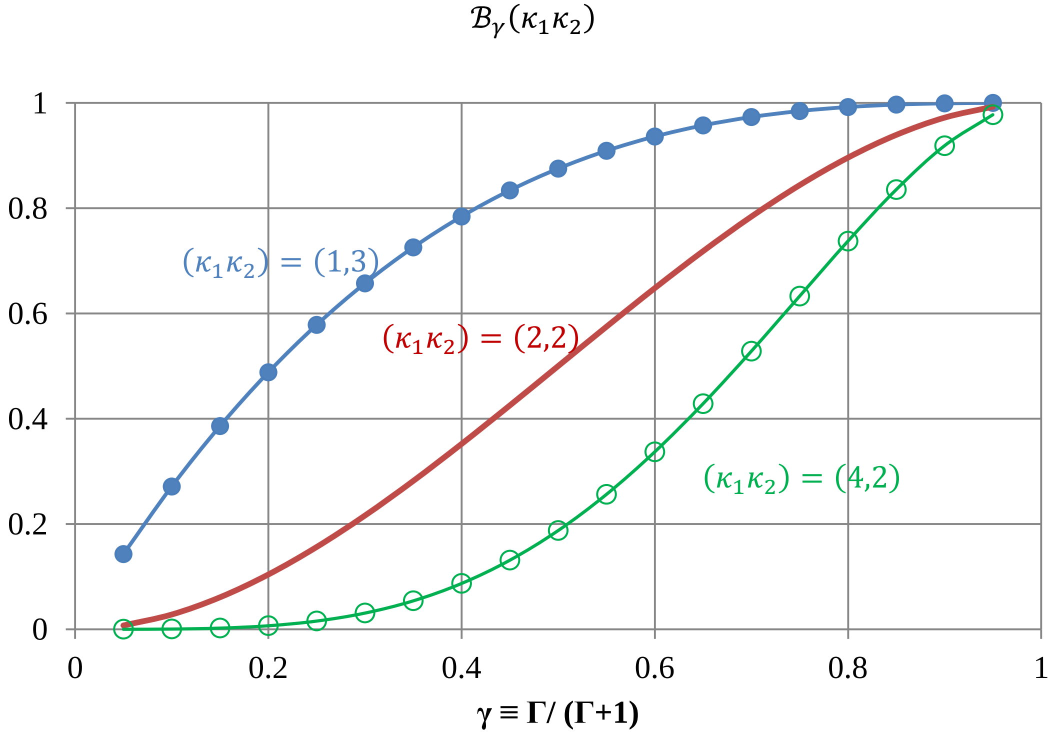

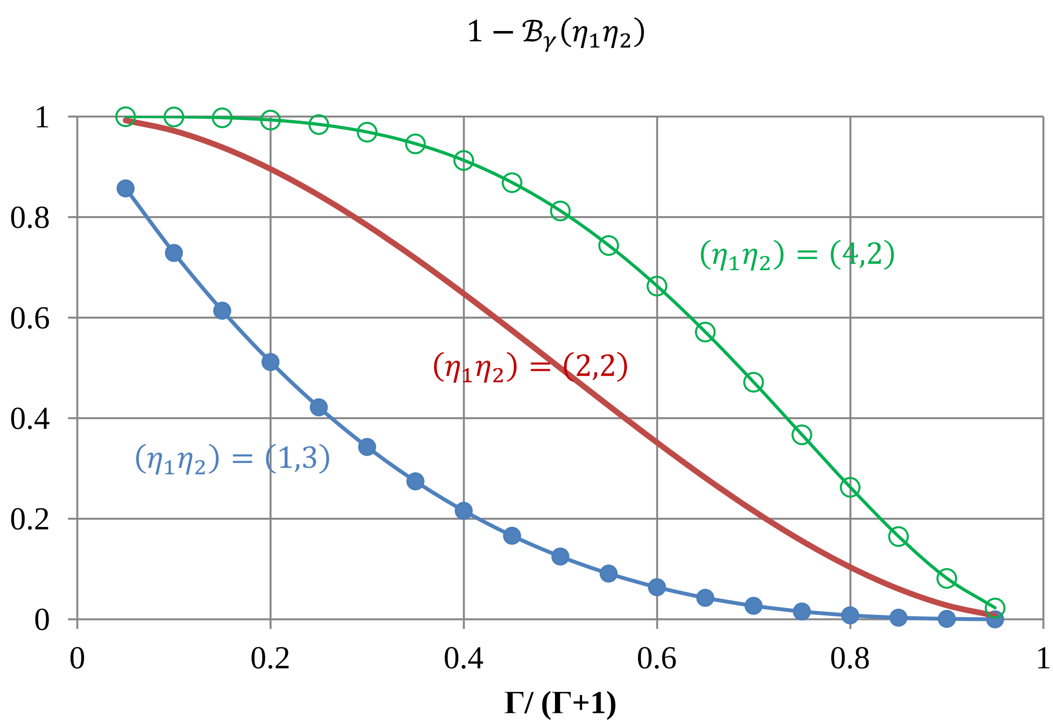

Figure 5: The distribution of employment growth in the data and task-based model

A. Data *B. Model*

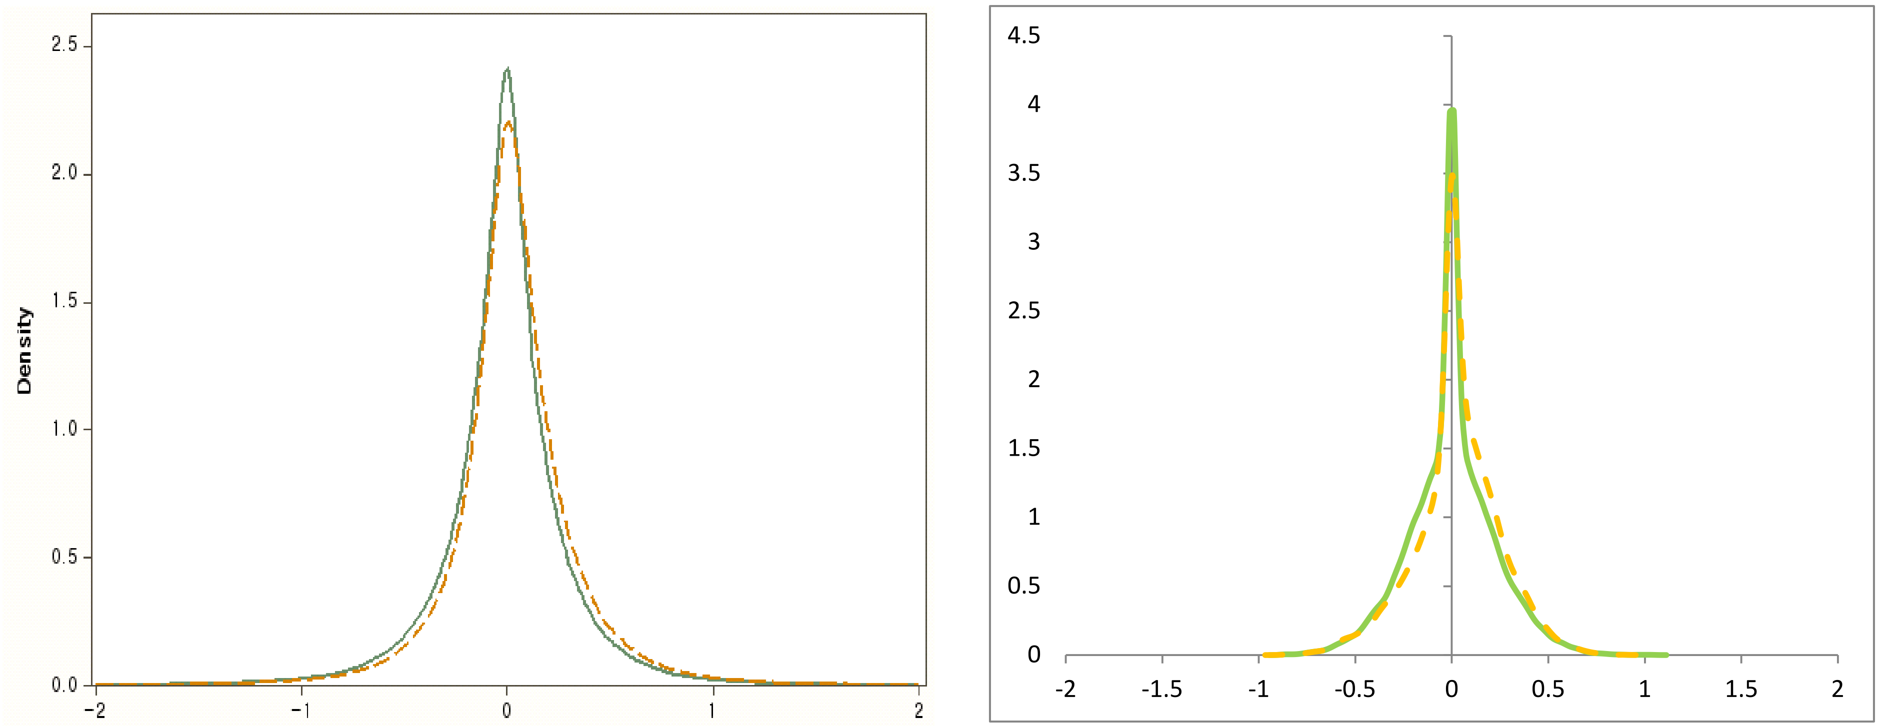

NOTE: The left panel shows the distribution of employment growth in Korean data. The solid (green) line is the kernsel-density estimate of the unconditional distribution, and the dashed (orange) line is the distribution conditional on an investment rate in excess of 10 percent. The right panel shows the distributions in the task-based model discussed in section 3.2.

Figure 6: The response of employment to an investment subsidy

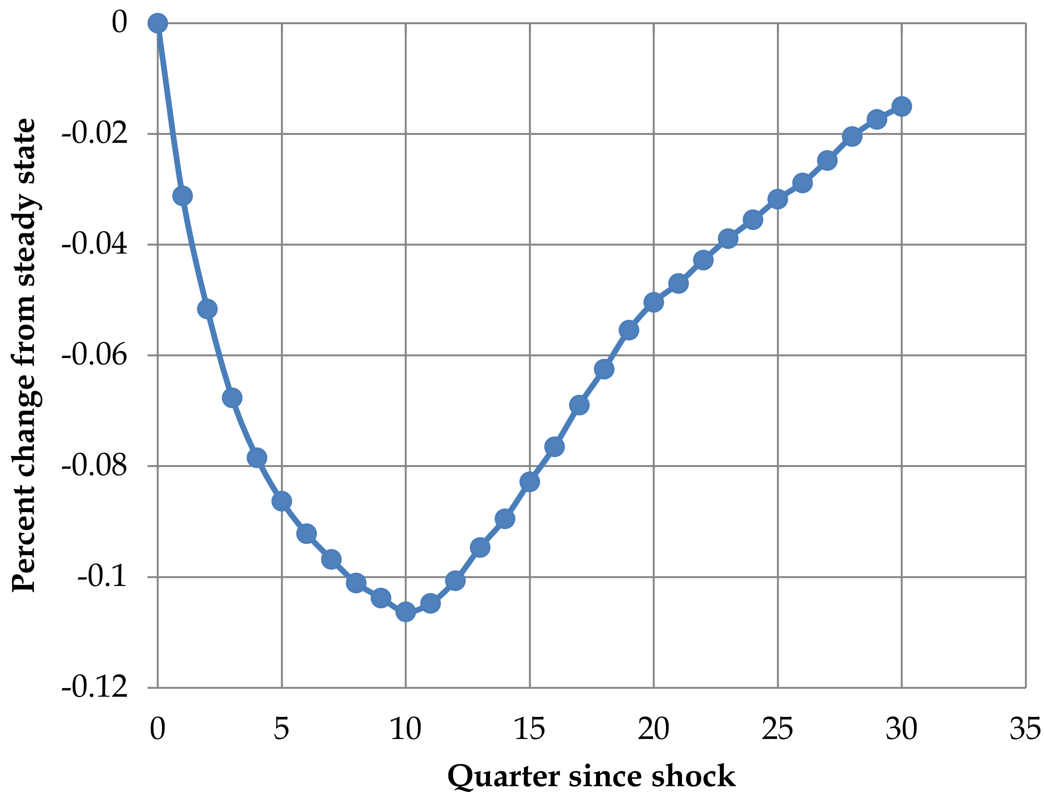

Table 1: Factor adjustment in Korea

Panel A: Frequency and magnitude of individual factor adjustments

Share of plant-year obs. for which $\frac{I}{K_{-1}} = 0$	0.47
Share of plant-year obs. for which $\Delta \ln n = 0$	0.135
Standard deviation of $\frac{I}{K_{-1}}$ (pooled across plants and time)	0.326
Standard deviation of $\Delta \ln n$ (pooled across plants and time)	0.266
Share of plant-year obs. with positive gross sales of machinery	0.143
Share of plant-year obs. with positive gross purchases *given* gross sales	0.746

Panel B: Joint factor adjustment

Share of plant-year obs. in which $\Delta \ln n < 0$, given $\frac{I}{K_{-1}} > 0$	0.395
Average decrease in $\ln n$, given $\Delta \ln n < 0$ and $\frac{I}{K_{-1}} > 0$	0.19
Average increase in $\ln n$, given $\Delta \ln n > 0$ and $\frac{I}{K_{-1}} > 0$	0.214

Table 2: The coincidence of positive investment and (net) negative employment growth

Panel A: By two-digit industry

	Share of plant-year obs. in which $\Delta \ln n < 0$, *given* $\frac{I}{K_{-1}} > 0$
Minimum among industries:	0.374 (“Other machinery”)
Median among industries:	0.419 (“Computer and office equipment”)
Maximum among industries:	0.458 (“Apparel”)

Panel B: By plant size

	Share of plant-year obs. in which $\Delta \ln n < 0$, *given* $\frac{I}{K_{-1}} > 0$
Plants w/ 10-24 workers:	0.358
Plants w/ 25-50 workers:	0.393
Plants w/ 51-99 workers:	0.412
Plants w/ 100+ workers:	0.476

Table 3: Baseline calibration

Parameter	Meaning	Value	Reason
β	Elasticity of output wrt n	0.50	Labor share
α	Elasticity of output wrt k	0.40	Capital share
ϱ	Discount factor	0.9875	Annual real interest rate = 5.2%
c_k^+	Purchase price of capital	1	Normalization
$-c_k^-$	Resale price of capital	0.923	Frequency of investment
c_n^+, c_n^-	Cost to hire & fire	33% of monthly wage	Frequency of net employment adj.
σ	Std. dev. of innovation to x	0.15	Avg. increase in $\ln n$, given $I/K_{-1} > 10\%$

Table 4: Empirical and simulated moments: The baseline model

Moment	Model $\delta_n = 0$	Model $\delta_n > 0$	Data
Prob($\Delta n = 0$):	0.148	0.006	0.135
Prob$\left(\frac{I}{K_{-1}} = 0\right)$:	0.468	0.438	0.47
Std dev of $\Delta \ln n$:	0.16	0.176	0.266
Prob$\left(\frac{I}{K_{-1}} < 0\right)$:	0.04	0.125	0.143
Share of qtly. obs. w/ $\Delta n < 0$ given $\frac{I}{K_{-1}} > 0$:	0	0.026	n.a.
Avg. qtly. decrease in $\ln n$ given $\frac{I}{K_{-1}} > 0$:	n.a.	-0.005	n.a.
Avg. qtly. increase in $\ln n$ given $\frac{I}{K_{-1}} > 0$:	0.114	0.108	n.a.
Share of yrly. obs. w/ $\Delta n < 0$ given $\frac{I}{K_{-1}} > 0.1$:	0.022	0.012	0.395
Avg. yrly. decrease in $\ln n$ given $\frac{I}{K_{-1}} > 0.1$:	-0.019	-0.016	-0.19
Avg. yrly. increase in $\ln n$ given $\frac{I}{K_{-1}} > 0.1$:	0.212	0.213	0.214

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 change in log employment; and the share of observations that involve disinvestment. 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. Of the two columns related to the model, the first pertains to the case with no worker attrition. The second assumes a quarterly attrition rate of 4%.

Table 5: Sensivitity analysis – non-technological theories

Moment	**Model** *Delivery lag*	*Disruption cost*	**Data**
Prob($\Delta n = 0$):	0.06	0.063	0.135
Prob$\left(\frac{I}{K_{-1}} = 0\right)$:	0.503	0.729	0.47
Std dev of $\Delta \log n$:	0.134	0.185	0.266
Prob$\left(\frac{I}{K_{-1}} < 0\right)$:	0.032	0.048	0.143
Share of qtly. obs. w/ $\Delta n < 0$ given $\frac{I}{K_{-1}} > 0$:	0.02 [orders]	0	n.a.
Avg. qtly. decrease in $\ln n$ given $\frac{I}{K_{-1}} > 0$:	-0.003 [orders]	n.a.	n.a.
Avg. qtly. increase in $\ln n$ given $\frac{I}{K_{-1}} > 0$:	0.059 [orders]	0.03	n.a.
Share of yrly. obs. w/ $\Delta n < 0$ given $\frac{I}{K_{-1}} > 0.1$:	0.125	0.108	0.395
Avg. yrly. decrease in $\ln n$ given $\frac{I}{K_{-1}} > 0.1$:	-0.035	-0.029	-0.19
Avg. yrly. increase in $\ln n$ given $\frac{I}{K_{-1}} > 0.1$:	0.168	0.306	0.214

NOTE: In regards to the definitions of the moments, see the Note to Table 4. The column labeled "Delivery lag" reports results for the model in which there is a quarterly probability of $\frac{1}{3}$ that a delivery of machinery is made. Note that, in the middle panel, we reported changes in labor conditional on positive *orders*. This illustrates that, at a quarterly frequency, positive orders and negative employment growth virtually never coincide. In the third panel, we report changes in labor conditional on positive *deliveries*, which corresponds to annual *measured* investment in the data. The column "Disruption cost" pertains to the model in which capital adjustment reduces revenue by τ percent. Both models are discussed in section 3.1.

Table 6: The maximum feasible values of φ in the model of labor-augmenting TFP

		s						
		0.4	**0.8**	**1.2**	**1.6**	**2**	**2.4**	**2.8**
	1.5	-0.25	0.375	0.5833	0.6875	0.75	0.7917	0.8214
	2	-1.5	-0.25	0.1667	0.375	0.5	0.5833	0.6429
	2.5	-2.75	-0.875	-0.25	0.0625	0.25	0.375	0.4643
ϵ	**3**	-4	-1.5	-0.667	-0.25	0	0.1667	0.2857
	3.5	-5.25	-2.125	-1.083	-0.563	-0.25	-0.042	0.1071
	4	-6.5	-2.75	-1.5	-0.875	-0.5	-0.25	-0.071
	4.5	-7.75	-3.375	-1.917	-1.188	-0.75	-0.458	-0.25

NOTE: For each pair (ϵ, s), the table reports values of $\hat{\varphi} = 1 - \frac{\epsilon - 1}{s}$. In the static version of the model with labor-augmenting technology, the elasticity of substitution between capital and labor, φ, must be less than $\hat{\varphi}$ in order for technology to be labor-saving. It follows that, since $0 \leq \varphi$, only the cells shaded in gray are feasible.

Table 7: Sensivitity analysis – factor-biased technical change

Moment	**Model** *Labor-aug TFP*	*Task-based*	**Data**
Prob($\Delta n = 0$):	0.155	0.161	0.135
Prob$\left(\frac{I}{K_{-1}} = 0\right)$:	0.50	0.47	0.47
Std dev of $\Delta \log n$:	0. 139	0.217	0.266
Prob$\left(\frac{I}{K_{-1}} < 0\right)$:	0. 0045	0.171	0.143
Share of qtly. obs. w/ $\Delta n < 0$ given $\frac{I}{K_{-1}} > 0$:	0	0.139	n.a.
Avg. qtly. decrease in $\ln n$ given $\frac{I}{K_{-1}} > 0$:	n.a.	-0.117	n.a.
Avg. qtly. increase in $\ln n$ given $\frac{I}{K_{-1}} > 0$:	0.078	0.093	n.a.
Share of yrly. obs. w/ $\Delta n < 0$ given $\frac{I}{K_{-1}} > 0.1$:	0.032	0.351	0.395
Avg. yrly. decrease in $\ln n$ given $\frac{I}{K_{-1}} > 0.1$:	-0.033	-0.185	-0.19
Avg. yrly. increase in $\ln n$ given $\frac{I}{K_{-1}} > 0.1$:	0.163	0.188	0.214

NOTE: In regards to the definitions of the moments, see the Note to Table 4. The column labeled "Labor-aug. TFP" reports results for the model in which production is given by a CES function with labor-augmenting technology. The column "Task-based" pertains to the model in which plant-level output is the aggregate over intermediate tasks. Both models are discussed in section 3.2.