The Joint Dynamics of Capital and Employment at the Plant Level

William B. Hawkins — Yeshiva University

Ryan Michaels — University of Rochester

Jiyoon Oh — Korea Development Institute

April 16, 2015

Abstract

This paper studies the joint dynamics of plant-level capital and labor adjustment using data from the Korean manufacturing sector. It highlights an under-appreciated fact: investment and reductions in employment frequently occur together. Employment declines occur in 36 percent of plant-year observations in which the investment rate exceeds 10 percent. These episodes have macroeconomic significance: they account for 35.5 percent of total capital accumulation. Viewed through the lens of a canonical model of plant-level dynamics in which factors are complements in production, these data are not easy to rationalize. They can be understood, however, within a more general framework in which capital can directly replace labor in certain tasks. This framework captures the intuitive notion that the elasticity of substitution between capital and labor should be higher in the long run than in the short run. The paper estimates this model and assesses its implications both in and out of sample.

Keywords: joint dynamics, investment, employment adjustment, elasticity of substitution, Le Chatelier principle, manufacturing sector, Korea.

JEL codes: D24, E22, E23, E24.

We are grateful to participants at several seminars and conferences, and especially to Daron Acemoglu, Omer Acikgoz, Mark Bils, Russell Cooper, John Haltiwanger, Matthias Kehrig, Ezra Oberfield, Devesh Raval, and Daniel Xu for helpful comments. All errors are our own. E-mail address for correspondence: `ryan.michaels@rochester.edu`.

1 Introduction

It is by now well known that manufacturing plants adjust employment only infrequently, but often by a large amount conditional on adjusting.[1] Similar statements are true for capital: the plant-level investment distribution puts substantial mass at and near zero, but the right tail of investment "spikes" represents a significant part of aggregate capital formation.[2] These twin observations have been used, separately for each factor, to learn about the nature of factor adjustment costs at the plant level. The resulting insights are important for understanding the determinants and dynamics of aggregate investment and potentially also the cyclical behavior of employment and unemployment.[3] However, much less attention, both empirical and theoretical, has been paid to how plants adjust both capital and labor *jointly* and to what this might imply for broader economic questions.[4] This paper seeks to fill this gap. We highlight a striking and under-appreciated fact about the joint behavior of investment and employment growth at the plant level, and explore what this joint behavior suggests about plant-level production technologies.

Our paper makes two main contributions. The first is empirical. We document the joint dynamics of plant-level capital and labor adjustment in data for the South Korean manufacturing sector during the period 1990-2006. Perhaps the most interesting finding relates to the behavior of employment at investing plants. In particular, reductions in employment are common among plants undertaking large investments, a fact that has not been emphasized to date. We find that 36 percent of investing plants reduce employment. Moreover, these episodes involve significant changes in both factors of production. The average *de*crease in employment among investing plants that reduce their workforce is nearly 20 log points, and almost as large as the average *in*crease in employment among investing

[1]See Davis and Haltiwanger (1990, 1992), Davis, Faberman, and Haltiwanger (2006, 2013), and Cooper, Haltiwanger, and Willis (2007).

[2]See Doms and Dunne (1994), Nilsen and Schiantarelli (2003), Cooper and Haltiwanger (2006), and Gourio and Kashyap (2007). The literature on investment (see also Doms and Dunne, 1998) has stressed the skewness and kurtosis of the investment distribution in addition to exact inaction.

[3]For investment, see Caballero, Engel, and Haltiwanger (1995), Caballero, Engel, and Haltiwanger (1997), Caballero and Engel (1999), Cooper, Haltiwanger, and Power (1999), Thomas (2002), Veracierto (2002), Gourio and Kashyap (2007), Khan and Thomas (2008, 2013), and Arellano, Bai, and Kehoe (2012); for the labor market, see Cooper, Haltiwanger, and Willis (2007), Veracierto (2009), Kaas and Kircher (2011), Elsby and Michaels (2013), Acemoglu and Hawkins (2014), Schaal (2012), and Moscarini and Postel-Vinay (2009, 2012).

[4]There are some exceptions. Theoretical analyses of models with multiple frictions include Bloom (2009), Reiter, Sveen, and Weinke (2012), and Bloom et al. (2012). Empirical analyses using plant-level data include Sakellaris (2001, 2004), Letterie, Pfann, and Polder (2004), Polder and Verick (2004), and Eslava et al. (2010). We discuss the latter papers in more detail in Section 3.2 below. There is also an empirical literature which uses sectoral-level data (Nadiri and Rosen, 1969; Shapiro, 1986; Rossana, 1990; Hall, 2004). Merz and Yashiv (2007) is an important recent contribution.

plants that add to their workforce. The investment undertaken at plants that are contracting employment is also substantial: total capital accumulation during such episodes accounts for more than one third of aggregate capital formation during our sample. Thus, these episodes are relevant to our understanding of the determinants of aggregate investment.

The second contribution of the paper is our analysis of the implications of these facts for theories of plant-level dynamics. Clearly, the co-occurrence of positive investment and employment declines is inconsistent with a canonical model of the literature, in which the plant-level production function is Cobb-Douglas in capital and labor (for example, Bloom, 2009).[5] In this case, whenever capital and labor are adjusted, both are adjusted in the *same* direction (Dixit, 1997; Eberly and van Mieghem, 1997). This reflects the fact that all productivity shocks are necessarily Hicks-neutral in this environment and so cannot induce the plant to substitute capital for labor. More generally, the empirical prevalence of investment occurring together with employment reduction is hard to replicate within any conventionally parameterized constant-elasticity-of-substitution (CES) production function, of which Cobb-Douglas is of course a special case.[6]

Much of our theoretical analysis concerns an alternative framework. We study a model in which the plant's output is an aggregate over a set of tasks, each of which can be performed by either capital or labor (Zeira, 1998; Acemoglu, 2010). In the *short run*, the plant responds to changes in business conditions by adjusting production of each task along the intensive margin—for example, by adjusting how much labor is assigned to each task (already) performed by labor. Changing the set of tasks performed by each factor requires a more fundamental reorganization of production, and so is done only from time to time. Such a reorganization entails a shift in tasks from one factor to another—for instance, machinery is deployed to take over tasks formerly performed by labor. The production function after this reorganization is the *long-run* production function. It exhibits a larger elasticity of substitution between capital and labor, and the more so the greater the shift in the allocation of tasks to factors.

The tasks-based framework offers a tractable and flexible framework for the analysis of plant-level factor adjustment dynamics. It allows for essentially arbitrary short- and long-run elasticities of

[5] Bloom (2009) analyzes both factors of production in a model with capital and labor adjustment costs. Other studies, such as Cooper, Haltiwanger, and Willis (2004) and Cooper and Haltiwanger (2006), consider a Cobb-Douglas function in both factors but concentrate out the flexible factor.

[6] We discuss less conventional parameterizations of CES production functions below.

substitution provided only that the second is larger than the first; that is, our framework provides a natural formalization of an old idea, the Le Chatelier principle of Samuelson (1947). Moreover, under an appropriate functional form assumption on how the comparative advantage of capital to labor varies across tasks, our framework allows for the substitution elasticity to be constant in the short run and also constant (but larger) in the long run. This facilitates our quantitative investigation, since the parsimony of the model implies that certain aspects of the joint dynamics are governed by a small number of parameters.

Our tasks-based model generalizes two classic models of production, and does so in ways that are appropriate to engage our plant-level data.[7] First, the "putty-clay" model of Johansen (1959) is closely related to the special case of our model in which the short-run production function is Leontief. The assumption of a fixed capital-labor ratio may be appropriate at the level of a single machine, but it seems restrictive as a model of the production technology of a manufacturing plant which uses many types of machines and labor and can substitute between them as economic conditions change. Second, the standard CES production function is the special case of our model in which the plant never adjusts the allocation of factors to tasks. Although there exist parameterizations of a standard CES model which can rationalize the co-occurrence of employment reductions with investment (for example, if the substitution elasticity is close enough to zero) we did not want to assume at the outset that all factor-augmenting technological progress must *always* have this effect in both the short and the long run. The general tasks-based framework allows, for instance, for an improvement in capital-augmenting technology to increase labor demand in the short run, but decrease it in the long run after the plant reorganizes its production process.

To assess whether our tasks-based model is consistent with the joint factor dynamics, we embed it in a dynamic stochastic setting and estimate it against our plant-level moments. Each plant in the model operates a production technology along the lines described, and faces a downward-sloping demand curve. The plant faces two driving forces: (factor-neutral) shifts in product demand, and innovations to factor-augmenting technology. The latter is what drives changes in the desired assignment of factors to tasks. However, adjustments along any margin—capital, employment, and

[7] Our model is also reminiscent of Houthakker (1955-56), Jones (2005), and Caselli and Coleman (2006), who derive production functions by aggregating over lower substitution-elasticity production "possibilities." In the same vein, our long-run plant-level production function is an envelope of the short-run production functions which arise for different factor-task allocations. See also Nakamura and Nakamura (2008), Antony (2009, 2014), Nakamura (2010), and Léon-Ledesma and Satchi (2011) for other, more reduced-form, approaches to time-varying substitution elasticities.

the assignment of factors to tasks—cannot be implemented freely. The costs of adjusting capital and labor help reproduce the marginal distributions of investment and employment growth, which show inaction: there are years when neither factor is changed. The friction in adjusting the assignment of factors to tasks allows for a distinction between short- and long-run elasticities of substitution.

The model is estimated by the method of simulated moments. We wish to remain agnostic about how technical progress should be incorporated in the production function, and so we consider, separately, capital- and labor-augmenting technology.[8] The tasks-based framework adds explanatory power in each case, but as we shall see, for different reasons.

When technology is capital-augmenting, the tasks-based model provides an excellent fit to the plant-level dynamics. In particular, it replicates very accurately salient moments of the marginal and joint distributions of capital and employment adjustments. The model accounts for episodes of opposite-direction factor adjustment as times when the plant changes the allocation of factors to tasks. Because of the tasks margin, the estimated model exhibits a large difference between the short- and long-run elasticities of substitution, although the latter of these is notably higher than conventional estimates (Chirinko, 2008; Raval, 2014; Oberfield and Raval, 2014).

Under labor-augmenting technology, the tasks margin proves to be more difficult to identify. This is because, even if the tasks margin is shuttered, labor-augmenting technical change can account reasonably well for episodes with positive investment and reductions in employment if there is a very low, near-Leontief elasticity of substitution. However, it is hard to see how the long-run elasticity of substitution could be zero. In this case, then, the tasks-based framework remains critical, but for a different reason: as we show, it enables this model to engage out-of-sample evidence on the long-run substitutability of factors.

We go on to generalize our tasks-based framework to include multiple labor inputs. The reason we do this relates to a dimension of the data we have not yet discussed: wages. When investment is positive and employment falls, we find that the average real wage increases, and by more than twice typical real wage growth. This suggests that the investment in these periods is likely accompanied by a shift in the composition of the workforce toward higher-wage workers who are more complementary to capital. Building on Goldin and Katz (1998), we show how to extend our framework in a tractable

[8] Diamond, McFadden, and Rodriguez (1978) is the classic statement on the challenges of identification for factor-augmenting CES production functions. Klump, McAdam, and Willman (2007) and Léon-Ledesma, McAdam, and Willman (2010) have made important strides in the context of static models.

way to replicate both moments on quantities (the joint adjustments of the factors) and wages.

Last, we pursue the idea that the failure of the canonical model of Cobb-Douglas production to generate episodes of positive investment and negative employment growth might arise not from an incorrect specification of the production function, but instead from other auxiliary assumptions. We modify our canonical model by allowing in turn for several potentially important features we abstracted from in the interest of tractability. These include time aggregation, delivery lags with respect to capital, disruption costs of installing machinery, and required replacement investment. Certain mechanisms can offer some explanatory power, but we argue that no extension provides a compelling quantitative account for the joint dynamics of factor adjustment. We view this as evidence in support of our premise that these dynamics are indeed informative about the production technology.

The remainder of the paper proceeds as follows. To fix ideas, Section 2 outlines a canonical model of dynamic factor demand under Cobb-Douglas technology and characterizes its predictions for joint factor adjustment. Section 3 describes our data and assesses the canonical model's predictions. Section 4 introduces the tasks-based production model and establishes its key properties concerning short- and long-run substitution elasticities. In Section 4, we also shed light on the joint dynamics of capital and labor implied by this production technology and contrast them with results in more standard CES frameworks. Section 5 describes the full dynamic model of tasks, which incorporates adjustment frictions on capital, employment, and the allocation of factors to tasks. Section 6 describes the parameter estimates for this model and evaluates its ability to replicate key features of the empirical factor adjustment distributions. This section also discusses the limitations of our identification strategy. Section 8 considers several alternative explanations of the plant-level dynamics that are rooted in the canonical Cobb-Douglas framework. Finally, Section 9 concludes.

2 A canonical framework

To organize our analysis of the data, it is instructive to first outline a simple theoretical framework. The model is due to Dixit (1997) and Eberly and van Mieghem (1997), and combines two influential approaches to capital and labor adjustment (Abel and Eberly, 1996; Bentolila and Bertola, 1990). The model is also part of the backbone of recent quantitative analyses of plant dynamics and their aggregate implications (e.g. Bloom, 2009). For these reasons, we refer to this framework as the *canonical model*. Since Dixit as well as Eberly and van Mieghem provide a lucid treatment of the problem, we focus on the key empirical predictions. We compare these to data in the next section.

Plants operate a Cobb-Douglas production function.[9] Physical output is given by

$$y = k^{\alpha} n^{1-\alpha}, \tag{1}$$

where the plant's stock of machinery is represented by k, the size of its workforce by n, and $0 < \alpha < 1$. The plant faces a downward-sloped product demand curve

$$y = \zeta p^{-\epsilon}. \tag{2}$$

Here ϵ is the (constant) elasticity of demand, and ζ represents a demand shifter. Combining (1) and (2) implies that the plant's revenue is

$$p(y)y = \zeta^{\frac{1}{\epsilon}} \left(k^{\alpha} n^{1-\alpha}\right)^{\frac{\epsilon-1}{\epsilon}}. \tag{3}$$

We assume that ζ obeys a geometric random walk. This is consistent with the spirit of studies, such as Foster, Haltiwanger, and Syverson (2008), that find that idiosyncratic shocks to plant revenue are very persistent.[10]

The factors of production, k and n, are subject to adjustment costs. In particular, with respect to employment, the plant pays c^{+} per hire (to recruit, train, and so on) and expends c^{-} per separation (which may be thought of as a statutory severance cost). With respect to capital, we assume plants sell used equipment at a discount. Thus, if the purchase price is normalized to unity, the resale price is assumed to be $p_s < 1$. More formally, where (k_{-1}, n_{-1}) are the initial levels of capital and

[9] The analysis applies to any revenue function that is monotonic in x; homogeneous of degree one in the tuple (k, n, ζ); and jointly concave and supermodular in (k, n, ζ). For concreteness, we use the Cobb-Douglas form.

[10] For simplicity, we assume that the only factor-neutral shock is to demand. Our analysis would be essentially unchanged if we allowed for a factor-neutral productivity shock in addition.

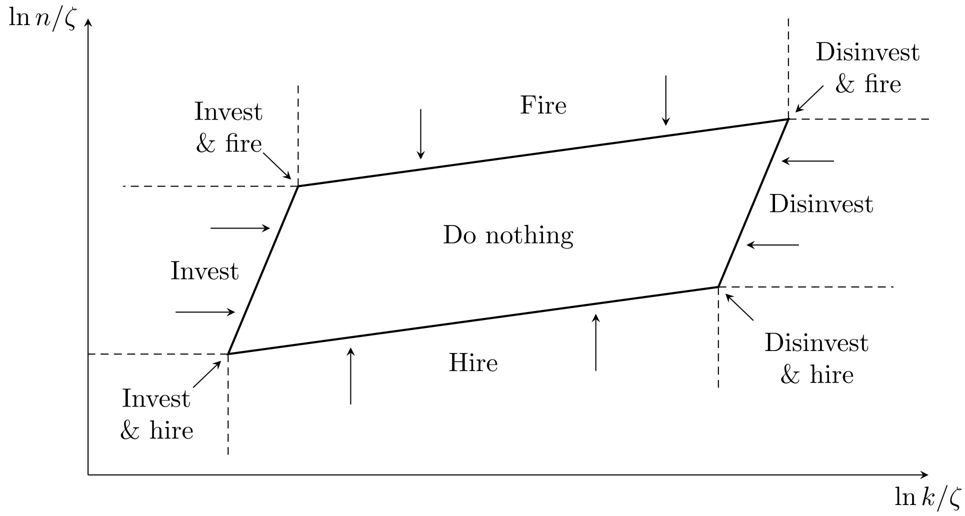

Figure 1. Factor demand policy in the canonical model.
Notes: This illustrates the optimal capital and labor demands in the canonical model. The horizontal axis is the level of capital normalized by profitability, ζ. The vertical axis is the level of employment normalized by ζ.

employment, the adjustment costs incurred by choosing inputs (k, n) are given by[11]

$$\mathcal{C}_n\left(n, n_{-1}\right)=\begin{cases} c^{+}\left(n-n_{-1}\right) & \text{if } n \geq n_{-1} \\ c^{-}\left(n_{-1}-n\right) & \text{if } n<n_{-1} \end{cases} ; \quad \mathcal{C}_k\left(k, k_{-1}\right)=\begin{cases} k-(1-\delta) k_{-1} & \text{if } k \geq(1-\delta) k_{-1} \\ p_s\left[k-(1-\delta) k_{-1}\right] & \text{if } k<(1-\delta) k_{-1} \end{cases} . \tag{4}$$

δ denotes the depreciation rate. (For simplicity, we follow Dixit in setting $\delta = 0$ in this section. We will allow for a positive depreciation rate in our numerical work in Section 6.) As is well known, the form of the adjustment frictions implies a wedge between the marginal costs of adjusting (any one factor) up versus down.[12] Therefore, there will be states of nature (ζ) in which it is not optimal to adjust; there is an optimal degree of inaction.

The plant chooses capital and labor each period in order to maximize the expected present discounted value of revenues, less expenditures on wages and adjustment costs.[13] As proved by Dixit (1997) and Eberly and van Mieghem (1997), the plant's decision rules take a form that is illustrated in Figure 1. The horizontal axis is a normalized measure of capital demand, $\ln k/\zeta$: if this is small,

[11]Throughout, a prime ($'$) indicates the next period value, and a subscript $_{-1}$ the prior period value.
[12]This abstracts from fixed costs of adjusting. Jovanovic and Stolyarov (2000) analyze a model where two factors, which are complements in production, are subject to fixed costs of adjustment. They show that adjustment is staggered across the two, but the plant does *not* reduce one when it increases the other.
[13]The model of Section 5 also encompasses the canonical model as a special case. Thus, the Bellman equation can be seen formally there.

capital is low relative to its profitability and so investment is more likely to be undertaken.[14] The vertical axis is $\ln n/\zeta$, a normalized measure of labor demand.

The two shallow-sloped boundaries of the parallelogram depict the hiring and firing first-order conditions, taking as given the choice of k/ζ. The gap between these barriers is the inaction region, where the plant sets $n = n_{-1}$. In this range, the marginal value of employment is neither high enough to warrant paying the hiring cost nor low enough to justify incurring the separation cost. If profitability increases substantially but the plant does not adjust, its $\ln n/\zeta$ will decline and (eventually) reach the lower barrier of the parallelogram; hiring is then optimal. The same idea applies (but in reverse, of course) with respect to firing. Note that the labor demand schedules are increasing in $\ln k/\zeta$ because of the supermodularity of the profit function: the marginal value of employment, whether hiring or firing, is increasing in capital.

The more steeply-sloped boundaries of the parallelogram in Figure 1 depict the investing (on the left) and disinvesting (on the right) first-order conditions, taking as given the choice of n/ζ.[15] We depict the inaction space in the capital dimension as relatively wide, in light of our results to follow (Section 3) that investment is less frequent than employment adjustment. The capital demand functions are upward-sloping, again due to supermodularity.

Figure 2 depicts the dynamics implied by these policy rules in a particular example. We suppose the plant begins inside the inaction region. It then faces a sequence of increases in ζ. The pair $(\ln k/\zeta, \ln n/\zeta)$ first drifts southwest along a 45^o line, as profitability increases relative to k and n.[16] Since the cost of hiring is relatively small, the hiring barrier is reached first. If ζ continues rising, the plant continues increasing employment to remain on its first-order condition—the lower barrier in the parallelogram—but, for the time being, it does not adjust capital. Accordingly, the pair $(\ln k/\zeta, \ln n/\zeta)$ moves leftward along the hiring barrier.

When the plant reaches the southwest corner, it invests and hires. The plant would remain

[14] The assumptions that the revenue function is homogeneous and ζ obeys a random walk imply that the present value of the firm scales linearly with ζ. Thus, the firm's problem can be expressed in terms of the normalized factor demands, k/ζ and n/ζ.

[15] The homogeneity and supermodularity of the firm's profit function imply that the capital demand schedules, written in $(\ln k/\zeta, \ln n/\zeta)$-space, have to be more steeply sloped. See Dixit (1997) and Eberly and van Mieghem (1997) for details.

[16] To be more precise, the pair $(\ln k/\zeta, \ln n/\zeta)$ moves along a 45^o line if the rates of attrition and depreciation are zero. If depreciation is positive, for instance, then capital falls by more if investment is not undertaken, so $(\ln k/\zeta, \ln n/\zeta)$ travels along a more shallow ray through the parallelogram. Our conclusions here—namely, the firm adjusts k less often and, when it does so, it adjusts n in the same direction—still apply as long as the gap between the attrition and depreciation rates is not large relative to the cost of adjusting capital. That is a safe assumption to impose.

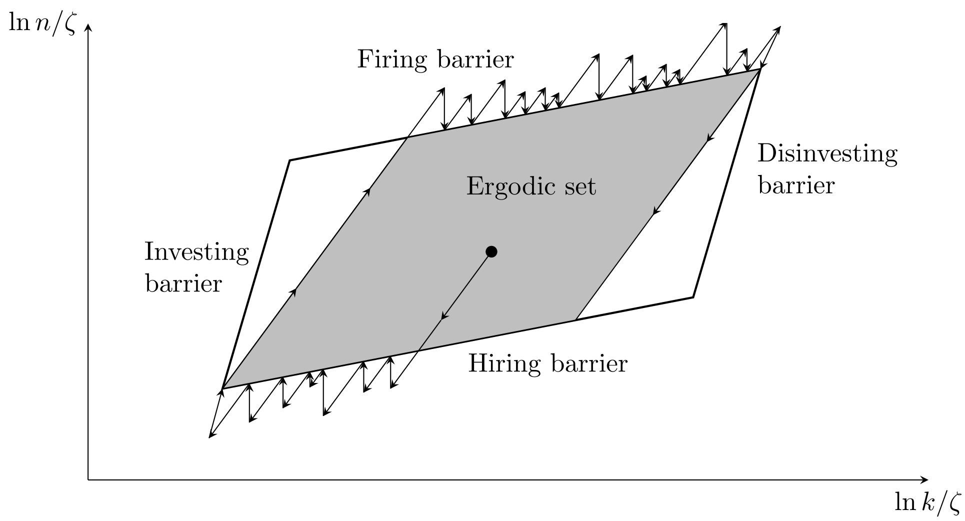

Figure 2. The joint dynamics of capital and employment in the canonical model.
Notes: This traces out the ergodic set of $(k/\zeta, n/\zeta)$ in the canonical model. In this example, the plant begins at the "dot" in the middle of the inaction region.

here if ζ continued to rise. If, on the other hand, ζ declines, the pair $(\ln k/\zeta, \ln n/\zeta)$ retreats inside the parallelogram. If ζ falls further, the plant will reach the firing barrier next, since the cost of reducing employment is smaller than the cost of disinvesting. As long as profitability continues to fall, the plant moves rightward along the firing barrier (since k/ζ is increasing). It will, at some point, reach the northeast corner, where it will fire and disinvest. When profitability does begin rising again, the pair $(\ln k/\zeta, \ln n/\zeta)$ retreats once more inside the parallelogram.

This pattern of adjustments traces out the ergodic set of the pair $(\ln k/\zeta, \ln n/\zeta)$, as highlighted by the shaded region in Figure 2. This set is bounded on the left by the 45° ray that extends northeast from the southwest corner of the parallelogram and bounded on the right by the 45° ray that extends southwest from the northeast corner. No matter the initial value of $(\ln k/\zeta, \ln n/\zeta)$, a plant will eventually enter this set. Once inside, the plant will never leave.

Critically, this set does not contain any points where the plant adjusts capital and employment in opposite directions. Thus, though the factors are adjusted infrequently, the complementarity of the two ensures that, *conditional on* adjusting, capital and employment move in the same direction. The empirical implication is that, *if one sees a plant invest, it should simultaneously hire.*

In the next section, we ask whether this striking prediction is supported in data. Before we proceed to that test, it is worth remarking that the canonical framework just presented abstracts from several features of the environment facing real-world manufacturing plants that could potentially be important for the joint dynamics of capital and labor. These include time aggregation, alternative models of wage and price determination, disruption costs of installing machinery, and required replacement investment. We return to the canonical framework and consider these issues in Section 8. However, the main implications of the model studied here will remain intact.

3 Facts about factor adjustment

This section consists of two parts. We first introduce our data and summarize a few facts about the marginal distributions of investment and employment growth at the plant level. We then take up the joint dynamics of factor adjustments, and test the key prediction of the canonical model.

3.1 Marginal distributions of capital and labor adjustment

To investigate the joint dynamics of capital and labor at the plant level, we use data from the Korean Annual Manufacturing Survey. The survey covers all manufacturing establishments with at least 10 workers and includes observations on the size of the plant's workforce and investment. The data are annual. Our analysis covers the period 1990-2006, which gives us 508,209 plant-year observations.[17] Our description of investment pertains to machinery investment specifically. Throughout, we define a plant's investment rate as i/k_{-1}, where i is real investment and k_{-1} is the real machinery stock at the end of the prior year.[18] Employment growth is constructed as the log change, $\Delta \ln n = \ln n - \ln n_{-1}$.

Before we turn to their joint dynamics, we begin with a brief description of the marginal distributions of net investment and employment growth. Figure 3 plots these distributions, and panel A of Table 1 provides some summary statistics.[19]

A first noteworthy feature of the data is that capital is adjusted less often than employment. Investment is reported to be exactly zero in almost 48 percent of all plant-year observations, whereas

[17] After 2006, plants with 5 to 10 workers were included in the survey. For comparability across time, we use data only through 2006.

[18] The dataset includes information on gross purchases and sales. What we refer to as real investment is the difference of the two, deflated by an equipment price index. See Web Appendix B for further details.

[19] Note that the bar at zero in Figure 3 exceeds 48 percent in height since it also includes some plants with values for investment which are very small, but not precisely zero.

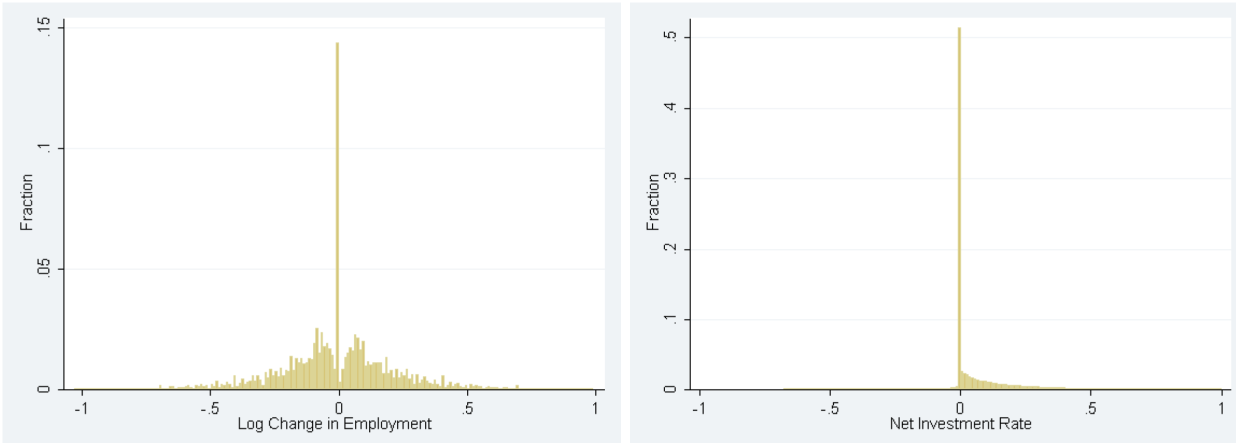

Figure 3. Factor adjustment distributions.

Notes: This plots the distributions of the log change in employment (top panel) and the net investment rate (bottom panel). The vertical axis is the fraction of plant-year observations in that bin.

employment growth is zero 14 percent of the time.[20] There is, of course, some heterogeneity across plants with respect to how often each factor is changed, but in about 80 percent of plants, it is adjusted more often than machinery. These facts imply that, viewed through the lens of the canonical model, capital is the more costly-to-adjust factor (that is, its (S, s) inaction band will be relatively wide).

Second, as illustrated in Figure 3, the basic shapes of the two distributions are very different. In particular, the distribution of investment is right-skewed, whereas the distribution of employment growth is roughly symmetric. More precisely, just 4.4 percent of investment observations are negative, and almost 48 percent of observations are (strictly) positive. In contrast, the share of observations with negative and negative employment growth are more similar (44.5 and 41.4 percent, respectively).

Third, if factors are adjusted, the size of the adjustments are substantial. A one-standard deviation change in employment amounts to over 22 log points, and the typical investment represents a gross increase in capital of about 50 percent. To make this point another way, the numbers in the table show that 47 percent (.225/.478) of positive investment episodes involve a gross increase in

[20]This frequency of zero investment is higher than in other datasets. Cooper and Haltiwanger (2006) report that 8 percent of plant-year observations in their U.S. data show zero investment. However, these results are derived from a balanced sample of large plants, which adjust more often. At plants in our Korean data with more than 100 workers, the frequency of zero investment is just 10.6 percent. Even these plants, though, adjust capital less often than labor. We also note that, although our sample includes plants that enter and/or exit over the period under study (it is unbalanced), we drop the years of entry and exit themselves (since our measures of factor adjustment are not well defined for such plants). Our results are virtually unchanged if we drop an additional year after entry and an additional year before exit.

Panel A: Single-factor adjustment		
Statistic	$\Delta \ln n$	i/k_{-1}
Fraction of plant-year obs. with…		
…no change	0.141	0.478
…positive change	0.414	0.478
…negative change	0.445	0.044
…change > 0.2	0.153	0.225
…change > 0.5	0.030	0.115
Standard deviation	0.220	0.505
Panel B: Joint adjustment		
$\text{Prob}[\Delta n < 0 \vert i/k_{-1} > 0.1]$	0.360	
$\mathbb{E}[\Delta \ln n \vert i/k_{-1} > 0.1, \Delta n < 0]$	-0.186	
$\mathbb{E}[\Delta \ln n \vert i/k_{-1} > 0.1, \Delta n > 0]$	0.208	
$\text{Prob}[\Delta n < 0 \vert i/k_{-1} > 0.2]$	0.341	
One-year persistence of $\Delta \ln n$ given $i/k_{-1} > 0.1$ and $\Delta n < 0$	0.700	
Share of aggregate investment due to $i/k_{-1} > 0.1$ and $\Delta n < 0$	0.355	

Table 1. Summary statistics

Notes: This reports several moments regarding the distributions of capital and labor adjustment in the Korean manufacturing sector, 1990-2006.

capital of at least 20 percent. Almost one-quarter (.115/.478) of these episodes record an increase of at least 50 percent. The counterparts among positive employment adjustment episodes are 37 percent, and 7 percent, respectively. Any model of factor adjustments will require substantial idiosyncratic profitability shocks to replicate these observations.

3.2 Joint distribution of capital and labor adjustment

Taking our cue from the analysis of Section 2, we investigate the distribution of log changes in employment conditional on positive investment. We condition on positive investment because our earlier results suggest that capital is the more costly-to-adjust factor at the vast majority of plants. Thus, the theory's sharpest prediction is that this conditional distribution should lie everywhere to the right of zero. When we take this to data, though, we wish to guard against the possibility that many positive investment episodes are mis-measured zeros. Therefore, we condition, more precisely, on an investment rate of at least 10 percent.

Figure 4 plots this conditional distribution and contrasts it with the unconditional distribution of log employment changes. The figure shows that, although the conditional distribution is shifted

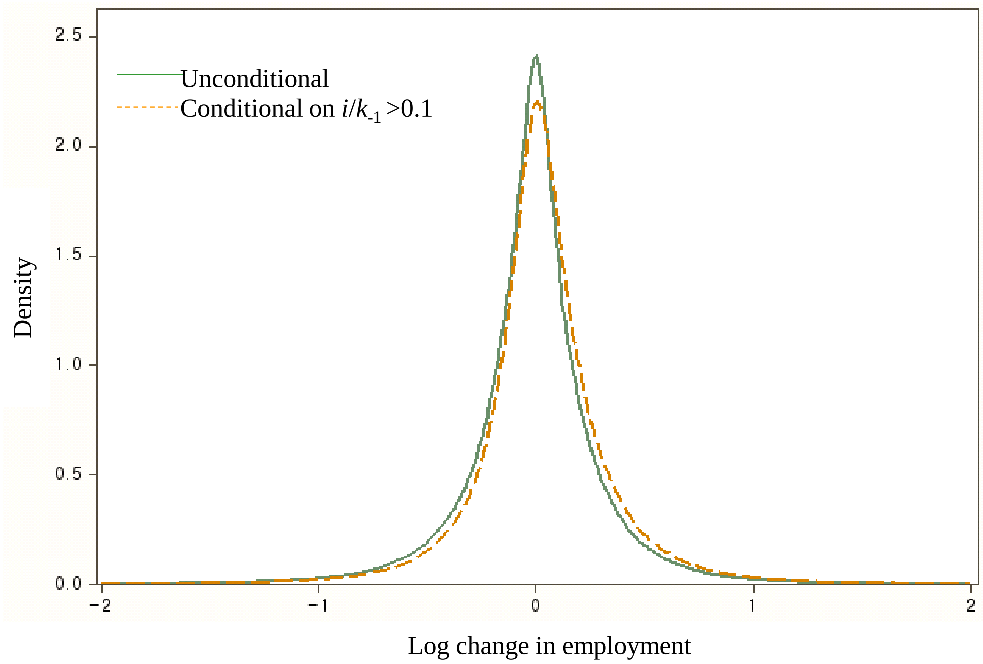

Figure 4. Employment adjustment conditional on investment: Korean data.
Notes: This plots the kernel density of the log change of employment at the plant-year level. The solid (green) line shows the unconditional distribution. The dashed orange line shows the distribution conditional on an investment rate in excess of 10 percent.

slightly to the right of its unconditional counterpart, the two distributions are remarkably similar. It seems that substantial employment declines often coincide with positive investment.[21]

Panel B of Table 1 summarizes the figure with a few key moments. First, among those plant-year observations that involve an investment rate greater than 10 percent, we compute the share in which employment declines. This is 36 percent: more than one-third of investment episodes (where $i/k_{-1} > 0.1$) involve a reduction in employment. Second, we calculate the average decline in employment among the subset of these plants that reduce their workforces when they invest. This is 18.6 log points. To put this in some perspective, the table also reports the average *expansion* in employment in years where plants both invest and hire (on net), which is 20.8 log points. Thus, among investing plants, the declines in employment at contracting plants are comparable to the increases in employment at expanding plants.

In addition, the employment losses in these investment episodes are not reversed especially

[21] We see only plants in our data, not the parent firms. Therefore, we do not know if workers separated from one plant are transferred to another plant within the firm. However, the vast majority of smaller plants are likely single-establishment firms, and the co-occurrence of positive investment and employment decline at small plants is comparable to that in the full sample.

quickly. To see this, we perform an exercise along the lines of Davis and Haltiwanger (1992). For each plant that sheds workers in year t, we compute the share of the decline that is present one year later. Aggregating across plants (and weighting by the plant's share of total year-t job destruction), we obtain a measure of the persistence of employment loss. We do this for all plants that report job destruction and for the sub-sample of plants which both reduce employment and invest. Our calculations show that, for both all contracting plants and for the sub-sample of investing plants, about 70 percent of the employment decline persists.

It is also important to emphasize that the amount of investment undertaken at times when employment falls is substantial. That is, opposite-direction adjustments of capital and labor do not (only) occur in periods when investment is small. There are two ways to illustrate this. First, 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 an investment spike in Cooper and Haltiwanger (2006). The probability of negative employment growth, conditional on a 20 percent threshold, remains high at 34.1 percent. Second, 35.5 percent of all investment in our data occurs in plant-years in which the investment rate exceeds 10 percent and employment growth is negative. Thus, these periods account for a significant portion of aggregate capital accumulation.

We have dug a little deeper to consider how these moments regarding the joint dynamics vary across a number of sub-samples. Results are reported in Web Appendix B. We find, in short, that the co-occurrence of investment with employment reductions appears to be a general phenomenon, not one restricted to particular types of plants, industries, or years.[22] For example, its prevalence varies across two-digit industries in a narrow range, from 0.328 (machinery) to 0.409 (food products). We also emphasize that the frequency of such episodes is roughly constant over time, at least with the exception of the Asian financial crisis years of 1997-98, when it is *more* frequent. This suggests that a model of factor adjustment dynamics needs to be able to account for this phenomenon in an aggregate steady state, rather than by appealing to aggregate or sectoral shocks. It does, however, also offer a hint that understanding joint factor dynamics may be important in understanding aggregate fluctuations, something that is beyond the scope of the current paper.

Our main results are for net investment, constructed as gross purchases less gross sales. However,

[22] There are limits to the extent we can disaggregate the data, though. For example, our dataset does not include information on whether a plant participates in international trade. For that reason, we cannot divide the sample into exporters and non-exporters.

our results are largely unaffected if we look at gross investment. This is largely because sales of used equipment are rare: gross sales are positive in only 14.0 percent of plant-year observations. It is, however, worth noting that plants which sell equipment usually (nearly 75 percent of the time) purchase new equipment at the same time.[23] This suggests equipment sales may be associated with upgrading or restructuring at the plant level. Viewed in this light, it is noteworthy that episodes of positive net investment are somewhat more likely to be associated with a simultaneous sale of capital if they occur along with an employment reduction. Nearly one-quarter of episodes in which $i/k_{-1} > 0.1$ and $\Delta \ln n < 0$ also exhibit positive gross sales of capital. We view this as suggestive evidence in favor of our tasks-based model, which delivers employment reductions together with investment when the plant restructures production. However, since sales of capital are rare, we focus henceforth on gross purchases less sales, and leave a study of upgrading for future work.

Last, we briefly relate our results to the literature. Though the negative co-occurrence of investment and employment growth seems generally under-appreciated, a few earlier papers did report it in other contexts. Sakellaris's (2001) analysis of the U.S. Annual Survey of Manufacturers finds that, in years of investment spikes (when the 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.[24] Letterie, Pfann, and Polder (2004) and Polder and Verick (2004) studied German and Dutch data and also observed that employment often declines when investment is positive.[25] However, relative to this work, our empirical analysis is more narrowly focused on these opposite-direction movements. To our knowledge, we are the first to call attention to the size and persistence of the employment losses in these episodes, and to highlight the pervasiveness of the phenomenon across industries, size classes, and time. In addition, we relate our empirical results to the predictions of the canonical theory and propose a rich, alternative framework in which to study this dimension of the joint dynamics.

Additional work on joint adjustment dynamics (for example, Eslava et al., 2010) emphasizes that labor and capital adjustment are *on average* positively correlated at the plant level. This

[23] Cooper and Haltiwanger (2000) observe that gross purchases and sales frequently coincide in U.S. data also.

[24] Sakellaris mentions the result only in the working paper version of his paper. The published paper (Sakellaris, 2004) omits it.

[25] 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 firing, plants ought to always disinvest. This prediction is clearly violated in their (and our) data.

is true in our data as well: the employment growth distribution among investing plants (dashed line in Figure 4) does lie to the right of the unconditional distribution (solid line). However, the canonical model implies a much stronger restriction, namely, that there are no opposite-direction factor adjustment episodes. For this reason, the prevalence of these episodes seems informative about the production process at the plant level.

4 A tasks-based model of production

Broadly speaking, there are two ways to react to the prevalence of episodes in which investment is accompanied by a reduction in employment. One interpretation is that these reflect the replacement of labor by capital. The alternative view is that this apparent substitution is a mirage: inference about the production technology based on factor dynamics could be confounded by complications, unmodeled in the canonical theory, that are unrelated to the plant's production function. In this paper, we adopt the first view (although we consider the merits of the second in Section 8 below).

Therefore, in this section we introduce a tasks-based model of production that we use to interpret our empirical results. Section 4.1 describes the microfoundations for our approach and establishes results regarding the elasticity of substitution between capital and labor. Our tasks-based production function is more flexible than a standard CES alternative: intuitively, it allows us to model the idea that the elasticity of substitution between capital and labor will be higher when the plant undergoes a reorganization of its production process than at other times. Section 4.2 then explains why this additional flexibility can help account for the factor dynamics observed in our data.

4.1 The production function in the short and long run

The production of the final good is a CES aggregate of a large number of intermediate inputs, or tasks, following Acemoglu (2010). Formally, we assume that the number of units of the final good is given by

$$y = \left(\int_{-d}^{1+d} y_i^{\frac{\varphi-1}{\varphi}} \, \mathrm{d}i \right)^{\frac{\varphi}{\varphi-1}} . \tag{5}$$

Here $d \geq 0$, y_i denotes the production level of task $i \in [-d, 1+d]$ and $\varphi > 0$ is the elasticity of substitution between any two tasks. The production technology for task i takes the form

$$y_i = \begin{cases} \xi_k a_i k_i & \text{if task } i \text{ is performed by capital,} \\ \xi_n b_i n_i & \text{if task } i \text{ is performed by labor.} \end{cases} \tag{6}$$

The parameters ξ_k and ξ_n scale the productivity of capital and labor on all of their respective tasks. After aggregating across tasks, we shall see that these terms play the role of exogenous shifts to factor-augmenting technology (see (7)). The term a_i measures the productivity of capital in task i, and b_i indexes labor productivity in this task.

Tasks with index $i \in [-d, 0]$ can only be performed by capital; thus, $a_i > 0$ and $b_i = 0$ for such tasks. Similarly, tasks $i \in [1, 1+d]$ can only be performed by labor, so that $a_i = 0$ and $b_i > 0$ for such tasks. Finally, tasks $i \in (0, 1)$ can be performed by either factor, so that for such tasks both a_i and b_i are strictly positive. Notice that we allow for $d = 0$, so there may not exist any tasks which cannot be performed by both factors.

We require that a_i/b_i be strictly decreasing on $(0, 1)$, so that the tasks are ordered by increasing comparative advantage of labor. This structure implies that there exists a threshold task, $x \in (0, 1)$, such that tasks in the interval $[-d, x)$ are performed by capital and tasks in the interval $(x, 1+d]$ are performed by labor.[26] To simplify the analysis, it is also convenient to assume that a_i/b_i is continuously differentiable and that $\lim_{i \to 0^+} a_i/b_i = +\infty$ and $\lim_{i \to 1^-} a_i/b_i = 0$.

We first solve for the plant's production function conditional on the threshold task, x. For any triple (k, n, x), the plant allocates its factor inputs across each factor's respective set of tasks to maximize plant-wide output. Thus, the plant's production function is defined by

$$F(k, n, x) = \max_{\{k_i, n_i\}} \left[\int_{-d}^{x} (a_i \xi_k k_i)^{\frac{\varphi-1}{\varphi}} \, \mathrm{d}i + \int_{x}^{1+d} (b_i \xi_n n_i)^{\frac{\varphi-1}{\varphi}} \, \mathrm{d}i \right]^{\frac{\varphi}{\varphi-1}},$$

[26] This is shown formally as part of the proof of Proposition 1; see Appendix A. As demonstrated more clearly there, the simplicity of this result arises from the assumption that capital is homogeneous: each unit of capital can perform task i as well as any other. Thus, the shadow value of capital is independent of the identity of the task. The same idea applies to labor. This is a simplification that preserves tractability.

subject to the resource constraints

$$\int_{-d}^{x} k_i \,\mathrm{d}i \leq k \qquad \text{and} \qquad \int_{x}^{1+d} n_i \,\mathrm{d}i \leq n.$$

Proposition 1 summarizes the solution.[27]

Proposition 1. *For any $k, n > 0$ and any $x \in [0,1]$,*

$$F(k,n,x) = \left[A\left(x\right)^{1/\varphi}\left(\xi_k k\right)^{\frac{\varphi-1}{\varphi}} + B\left(x\right)^{1/\varphi}\left(\xi_n n\right)^{\frac{\varphi-1}{\varphi}}\right]^{\frac{\varphi}{\varphi-1}}, \tag{7}$$

where $A(x) = \int_{-d}^{x} a_i^{\varphi-1}\,\mathrm{d}i$ and $B(x) = \int_{x}^{1+d} b_i^{\varphi-1}\,\mathrm{d}i$.

We call (7) the *short-run* production function, as it takes x as given. It exhibits a constant elasticity of substitution between capital and labor equal to the underlying elasticity of substitution across tasks, φ. The threshold task, x, enters (7), in effect, like a component of factor-augmenting productivity. The extension of a given stock of capital to more tasks—that is, an increase in x—represents an increase in capital's productivity. This is because (5) displays diminishing returns to each task, so that applying capital to new tasks raises its marginal and average product. However, unlike conventional specifications of factor-augmenting technology, an increase in the productivity of capital via an increase in x necessarily, and simultaneously, *reduces* the productivity of the given stock of employment, n, since the latter is now spread over fewer tasks.

Next, we consider how the plant allocates factors to tasks, summarized by the threshold task, x. An increase in x, for example, represents the reallocation of some tasks previously performed by labor so that they are now performed by capital. While the plant's desire to increase x will be occasioned in our full model by an increase in its capital input, relative to labor, or by a change in capital- or labor-augmenting technology ξ_k or ξ_n, it is instructive first to consider the effect of changes in x *alone*. These changes enable the plant to achieve a higher elasticity of substitution than φ. As we argued in the Introduction, this added flexibility seems realistic: plants do have the ability to vary their capital intensity. At the same time, these adjustments may not happen frequently, but, rather, only in the long(er) run.[28]

[27] The Proposition's proof is included in Appendix A, like those of other Propositions in this section. Acemoglu (2010) solves a very similar problem in the context of an aggregate growth model.

[28] It is worth contrasting our approach with models such as Acemoglu and Zilibotti (2001), Acemoglu (2010), or

The *long-run* production function is defined as the maximum output that the plant can obtain if it is allowed to freely choose x, that is,

$$F^*(k,n) = \sup_{x\in[-d,1+d]} F(k,n,x).$$

With this notation, we can formalize the notion that the possibility of adjusting x increases the ability of the plant to substitute between factor inputs k and n (and the more so the larger the adjustment in x), as follows.

Proposition 2. *The elasticity of substitution between capital and labor for the long-run production function exceeds φ:*

$$-\left[\frac{d\ln\frac{F_k^*(k,n)}{F_n^*(k,n)}}{d\ln\frac{k}{n}}\right]^{-1} = \frac{\varphi}{1-\frac{\partial\ln\frac{A(x)}{B(x)}}{\partial\ln\frac{k}{n}}} \geq \varphi. \tag{8}$$

The intuition for Proposition 2 is straightforward. Consider an increase in the plant's capital input k (holding labor n fixed). If the threshold task x is fixed, the additional capital must be allocated only to tasks already undertaken with capital. This diminishes the marginal product of capital relative to the marginal product of labor. However, if x may be adjusted, the plant will also respond to the increase in capital by enlarging the set of tasks performed by that factor, that is, by increasing x. This increases $A(x)$ and decreases $B(x)$, thereby increasing the marginal product of capital and decreasing the marginal product of labor. Thus, the marginal products of the two factors change by less than when x is fixed: the elasticity of substitution is higher.

How much higher? The answer depends on how quickly the threshold task x optimally alters as k/n increases. This, in turn, depends on the slope of the comparative advantage schedule a_i/b_i near the threshold task. At one extreme, if a_i/b_i declines rapidly as x increases, then the plant will not alter x much even if it can freely do so: the long-run elasticity will not exceed the short-run elasticity by much. At the other extreme, if a_i/b_i is nearly constant, then x will be very responsive to changes in k/n, and the long-run elasticity will substantially exceed the short-run elasticity. This effect is captured by the presence of the factor $\{1-\partial\ln[A(x)/B(x)]/\partial\ln(k/n)\}^{-1} > 1$ in (8).

Acemoglu and Autor (2011), in which plants *always* choose the threshold task to maximize output. That is appropriate when studying steady-state or long-run questions exclusively. Because we are interested in short-run factor adjustment dynamics, we will allow (see Section 5 for more) for frictions in the process by which the plant adjusts the allocation of factors to tasks. This will endogenize this distinction between short- and long-run elasticities.

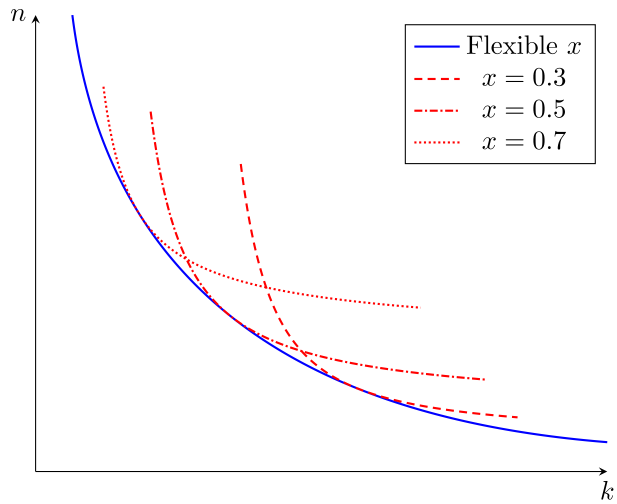

Figure 5. Short- and long-run production functions.
Notes: This shows the isoquants for the long-run (solid blue line) and short-run (dashed red lines) production functions. The three short-run isoquants correspond to three different values of the threshold task, x.

Proposition 2 can also be understood graphically. Figure 5 shows four isoquants corresponding to the same level of output. The solid isoquant is associated with the long-run production function $F^*(\cdot)$. The three dashed curves correspond to short-run production functions, $F(\cdot, x)$, for three different values of the threshold task x. Note that as k becomes more abundant, a plant will increase x in the long run, which in turn raises the marginal product of capital relative to labor. This means employment has to be reduced by more to remain on the same isoquant. In other words, the long-run isoquant will be less curved than the short-run isoquants that are tangent to it.[29]

We now show that, in a useful special case, our framework yields that not only the short-run, but also the *long-run* production function exhibits a constant elasticity of substitution. The key assumption is that the comparative advantage schedules a_i and b_i are power functions over the range $0 \leq i \leq 1$.[30]

Proposition 3. *Let $\psi > \varphi$, with $\psi \neq 1$. Define $\Gamma = (\psi - \varphi)/[\varphi(\psi - 1)]$.*

1. Suppose $\psi > \max\{\varphi, 1\}$. Let $\chi = [(\psi - \varphi)/(\psi - 1)]^{\frac{1}{\varphi - 1}}$, $d = 0$, $a_i = \chi i^{-1/(\psi-1)}$, and

[29] Figure 5 also makes clear how the firm's long-run production function $F^*(\cdot)$ is the envelope of the short run production functions $(k, n) \mapsto F(k, n, x)$ as x varies. As noted in the Introduction, this is reminiscent of Houthakker (1955-56), Jones (2005), and Caselli and Coleman (2006).

[30] Our quantitative analysis will rely mainly on the first part of Proposition 3, where $d = 0$ and the production structure simplifies somewhat. Thus, the reader may focus on this case without too much loss of generality.

$b_i = \chi(1-i)^{-1/(\psi-1)}$ *for all* $i \in (0,1)$*. Then the short-run production function is*

$$F(k,n,x) = \left[x^{\Gamma}\left(\xi_k k\right)^{\frac{\varphi-1}{\varphi}} + (1-x)^{\Gamma}\left(\xi_n n\right)^{\frac{\varphi-1}{\varphi}}\right]^{\frac{\varphi}{\varphi-1}}. \tag{9}$$

2. *Suppose* $\psi \in (\varphi, 1)$*. Let* $\chi = [(\psi-\varphi)/(1-\psi)]^{\frac{1}{\varphi-1}}$*,* $d=1$*, and*

$$a_i = \begin{cases} 1 & -1 \le i \le 0; \\ \chi(1-i)^{\frac{1}{1-\psi}} & 0 < i < 1; \\ 0 & 1 \le i \le 2, \end{cases} \quad \text{and} \quad b_i = \begin{cases} 0 & -1 \le i \le 0; \\ \chi i^{\frac{1}{1-\psi}} & 0 < i < 1; \\ 1 & 1 \le i \le 2. \end{cases}$$

Then the short-run production function is

$$F(k,n,x) = \left[(1-x)^{\Gamma}\left(\xi_k k\right)^{\frac{\varphi-1}{\varphi}} + x^{\Gamma}\left(\xi_n n\right)^{\frac{\varphi-1}{\varphi}}\right]^{\frac{\varphi}{\varphi-1}}. \tag{10}$$

3. *In both cases, the long-run production function is*

$$F^*(k,n) = \left[\left(\xi_k k\right)^{\frac{\psi-1}{\psi}} + \left(\xi_n n\right)^{\frac{\psi-1}{\psi}}\right]^{\frac{\psi}{\psi-1}}. \tag{11}$$

Proposition 3 is consistent with Proposition 2, in that the long-run elasticity must exceed the short-run elasticity, that is, $\psi > \varphi$. Proposition 2 also showed that, in order to generate a larger long-run elasticity, the threshold task must respond elastically to a change in the capital-labor ratio. Proposition 3 is consistent with this: the higher ψ, the slower the comparative advantage of capital (labor) decreases as the index of the task increases toward $i=1$ (decreases toward $i=0$) and thus the more elastically capital responds to a higher x.

In addition, under the functional form assumptions used in Proposition 3, we have two additional, sharp results. First, generating a long-run elasticity below 1 requires the existence of tasks which cannot be substituted between the factors at all, while for $\psi > 1$, this is not necessary.[31] Second, the short- and long-run elasticities φ and ψ are determined by separate features of the production technology. The first is equal to the elasticity of substitution across tasks; while the second is

[31] A converse result to Proposition 3 can also be established: if the short- and long-run elasticities of substitution are both constant, then the productivities of the factors must vary across tasks in this way, at least up to a monotone transformation of the interval of tasks $[-d, 1+d]$. Details are available on request.

determined by the slope of the comparative advantage schedule. This will be convenient in our numerical work, since we will be able to parameterize the two elasticities independently.

4.2 Joint factor dynamics

In this section, we illustrate why the tasks-based framework can account for opposite-signed movements in capital and labor. To isolate the intuition, we study the plant's response to a one-time, unanticipated, permanent shock to the production process. The assumption that the plant does not anticipate any future shocks allows us to characterize analytically what kinds of dynamics can occur. Section 5 allows for repeated shocks to productivity and demand.

We begin by studying the case in which the allocation of factors to tasks, summarized by the threshold task x, is fixed. Our results for this case would also apply to any other CES production function. We discuss the effect of changing x at the end of this section.

Formally, the fixed-x model is closely related to the canonical framework of Section 2. The plant faces the demand curve (2), and adjustment costs take the form (4). However, the Cobb-Douglas production function (1) is replaced by its tasks-based counterpart $F(k, n, x)$, as defined in (9) or (10). We assume that the demand shifter ζ and the two factor-augmenting technology shifters ξ_k and ξ_n have been constant for an indefinitely long time. In the current period, there is a one-time, unanticipated, permanent change to one of these. Following the shock, the plant chooses capital k and labor n in order to maximize the present value of profits, discounted with factor ϱ,

$$\max_{k,n} \left\{ \frac{1}{1-\varrho} \zeta^{\frac{1}{\epsilon}} \left[F(k,n,x)^{\frac{\epsilon-1}{\epsilon}} - wn \right] - \mathcal{C}_k\left(k, k_{-1}\right) - \mathcal{C}_n\left(n, n_{-1}\right) \right\}. \tag{12}$$

The form of (12) assumes that the plant will not adjust (k, n) again in the future; this is optimal since its environment will not change again. Since this problem shares the adjustment cost structure of the model in Section 2, it follows that the factor demand policies (expressed as k/ζ and n/ζ) can be summarized by the parallelogram in Figure 1 for any given ξ_k, ξ_n, and x.

We can now characterize the possible types of factor adjustment. For the same reason as in Section 2 a change in the demand shifter, ζ, will not lead capital and labor to adjust in opposite directions.[32] More interesting is the case of capital- or labor-augmenting technology, ξ_k or ξ_n; a

[32]We emphasize that the same would be true for any other factor-neutral shock.

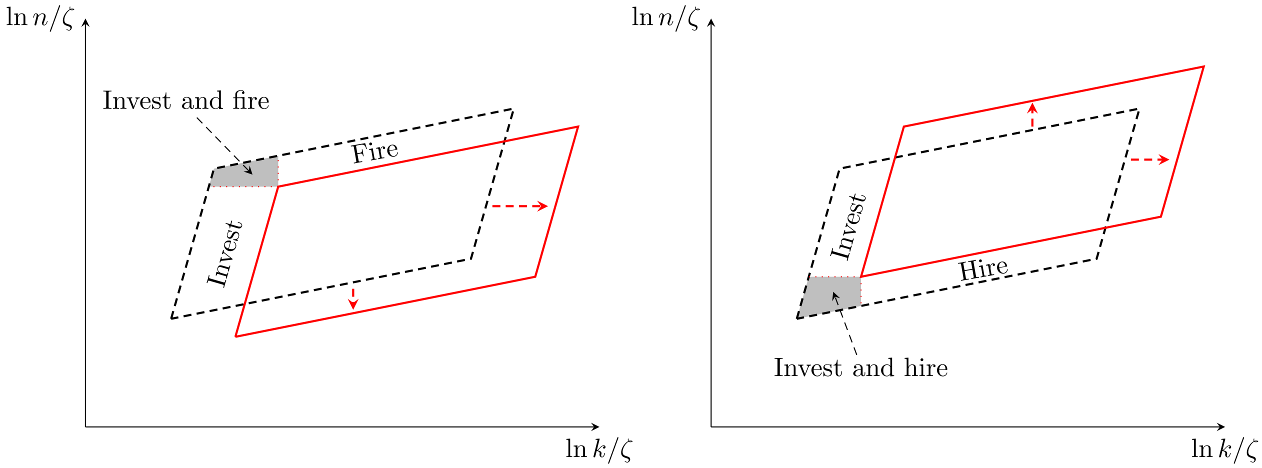

A. **B.**

Figure 6. Factor adjustment in response to an unanticipated permanent shock.
Notes: This illustrates how the capital and labor adjustments which occur in response to an unanticipated permanent shock depend on how the boundaries of the inaction parallelogram shift. See the text for more details.

shock to either of these induces a *shift* of the inaction parallelogram. The key observation for our purposes is that, if the shock is to induce the plant to invest and reduce employment, then some part of the original inaction region must lie above and to the left of the top left corner of the shifted parallelogram. When the boundaries of the parallelogram are upward-sloping, as shown in Figure 6, this occurs when the investing barrier shifts to the right and the firing barrier shifts down, as in panel A of Figure 6; firms with (k_{-1}, n_{-1}) in the shaded area are the ones which increase capital and decrease labor. (Notice that (k_{-1}, n_{-1}) must have been located within the original (dashed) inaction region, or else the plant would have adjusted its factor demands in a previous period.) Panel B, on the other hand, shows a different case in which the firing barrier shifts up rather than down. In this case, no part of the original parallelogram is associated with investment and employment reduction.

We now consider in turn the effects of changes to capital- and labor-augmenting technology.

Capital-augmenting technology. The movements of the firing and investment barriers in response to a change in capital-augmenting technology, ξ_k, depend on the relative strengths of two effects. The first is the substitution effect. If the two factors are substitutes ($\varphi > 1$), then for a given level of output, an improvement in ξ_k makes the cost-minimizing input mix more capital-intensive. This shifts the investment barrier to the right and the firing barrier down. The second effect is the scale effect, which arises because the plant faces a downward-sloping product demand curve.

Any improvement in technology reduces marginal cost, and leads the plant to produce more output by increasing both capital and labor inputs. This moves the investment barrier to the right and the firing barrier up. Thus, capital and labor may be adjusted in opposite directions only if the substitution effect is strong relative to the scale effect. The precise condition is that[33]

$$\varphi > \epsilon, \tag{13}$$

which is intuitive: the higher φ, the stronger the substitution effect, while the higher ϵ, the stronger the scale effect.

When the two factors are complements ($\varphi < 1$), on the other hand, it is a *decrease* in ξ_k which makes the plant raise its desired capital-labor ratio. When ξ_k decreases, the scale effect moves the investment barrier left and the firing barrier down, while the substitution effect shifts the investment barrier right and the firing barrier down. Thus, again, we can observe positive investment and a reduction in employment only if the substitution effect is strong enough; more precisely, if

$$\frac{\varphi}{1-\varphi} < \frac{1-s}{s}\frac{\epsilon}{\epsilon-1}, \tag{14}$$

where $s = rk/(rk + wn)$ denotes the capital cost share of the plant's factor payments (and $r \equiv 1 - \varrho$ is the user cost of capital). This condition is more likely to hold the lower φ (that is, the stronger the substitution effect) and the lower $\epsilon > 1$ (that is, the weaker the scale effect); it is also more likely to hold the larger s (that is, the larger the change in the marginal product of labor driven by a given change in ξ_k).[34]

To gauge the range of φ consistent with (13) and (14), suppose that $\epsilon = 4$, which is representative of the literature (see the references in Nakamura and Steinsson, 2008), and set $s = 0.5$, which holds approximately in our data. Then (13) holds for $\varphi > 4$, while (14) holds for $\varphi < 4/7 \approx 0.57$.[35]

[33] Proofs of the assertions about comparative statics in this section can be found in Web Appendix C.

[34] Note that the parameter condition that applies here in the case where k and n are gross complements is not quite the same as the one which would apply in the absence of adjustment frictions on capital and labor; an increase in ξ_k causes the frictionless choices of k and n to move in opposite directions if either $\varphi > \epsilon$ or $\varphi < 1 - (\epsilon - 1)s/(1 - s)$. The reason for the difference is that, for example, whether the firing barrier shifts up or down in response to an increase in ξ_k depends on whether labor demand increases holding k fixed; in the frictionless case, k is not fixed.

[35] The product demand elasticity measured by ϵ is the one that applies to *idiosyncratic* changes in plant output. The elasticity of substitution across industries might be (much) lower than 4. In that case, the co-occurrence of $i/k_{-1} > 0$ and $\Delta \ln n < 0$ could reflect a common shock that drives all plants up a steeply sloped industry product demand curve. However, the vast majority of variation in factor inputs in our data is within industry (e.g., plant-specific).

Labor-augmenting technology. We can similarly analyze the effects of a shift in labor-augmenting technology, ξ_n. Investment can again occur together with a reduction in employment under two scenarios. If capital and labor are substitutes, the required condition is that $\varphi > \epsilon$, that is, the same condition as (13). If the factors are complements, on the other hand, then we need that

$$\frac{\varphi}{1-\varphi} < \frac{s}{1-s}\frac{\epsilon}{\epsilon-1}, \tag{15}$$

which differs from (14) only in that the roles of the capital and labor shares, s and $1-s$, are switched. If $\epsilon = 4$ and $s = 0.5$, labor-augmenting technology can drive opposite-direction factor adjustment for $\varphi > 4$ or $\varphi < 0.57$, the same condition as in the capital-augmenting case.

Discussion. Before we turn to the general tasks-based framework, two comments are required. First, if φ is such that (13), (14), or (15) only *just* hold, then the model will generate some episodes of investment and employment reduction, but these will not involve the large *joint* movements in the two factors seen in our data. Thus, a quantitative account of joint factor dynamics under a fixed-x specification will require φ substantially further from unity than is suggested by the preceding analysis.

Second, we abstracted from depreciation. Intuitively, the need to replace capital lost to depreciation increases the likelihood that the plant will invest, and therefore also expands somewhat the range of parameters for which we can observe investment together with employment reduction. However, the main message of this section—that the elasticity of substitution φ needs to be sufficiently far from 1—is robust. We reintroduce depreciation in the full model in Section 5 and discuss how it would affect the results of this section in Web Appendix C.2.

Task allocation. Finally, consider a change in the distinctive feature of our tasks-based model, the allocation of factors to tasks, summarized by x. Although in our full model, x will be chosen endogenously by the firm, it is useful first to consider the response of factor demands to exogenous shifts in x.[36] Unlike a shift in ξ_k or ξ_n, an increase in x simultaneously amplifies the productivity of capital and degrades the productivity of labor, since tasks are reallocated toward k (and from n). When x does not begin too far from the optimum (as characterized in Proposition 3),

This suggests a more dominant role for idiosyncratic shocks.

[36] However, see Acemoglu (2010) and Acemoglu and Restrepo (2014) for models in which the set of tasks which may be performed by labor or capital may be exogenous from the plant's perspective.

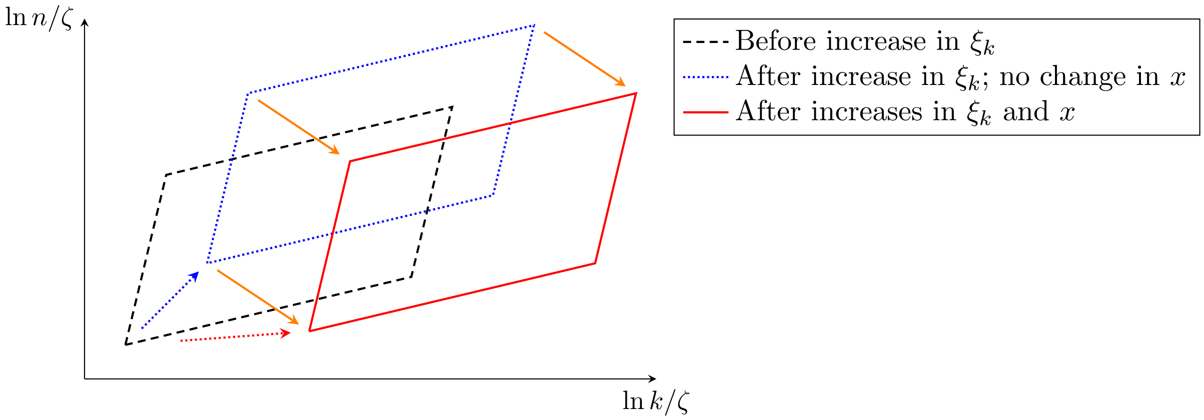

Figure 7. An increase in the threshold task, x, following an improvement in capital-augmenting technology ξ_k, shifts the inaction parallelogram downwards and to the right. See the text for more details.

this has a very stark implication for the movement of the barriers of the inaction parallelogram from Figure 1: they *always* move in the direction consistent with opposite-direction factor adjustment, as in panel A of Figure 6.[37]

Proposition 4. *For x in a neighborhood of x^*, an exogenous increase in x causes the barriers for capital to shift right and those for labor to shift down.*

A noteworthy feature of Proposition 4 is that it implies that an increase in x will induce positive investment along with a reduction of employment for any *short-run* elasticity of substitution φ. For example, the result is true even if φ is close to 1, a scenario in which changes in factor-augmenting technology alone will not induce opposite-direction movements in capital and labor.

Moreover, the Proposition also applies for any *long-run* elasticity of substitution $\psi > \varphi$. In particular, it applies even if $\psi < \epsilon$.[38] This is a particularly noteworthy result. Note that, if x were perfectly flexible, the model would not yield opposite-signed movements in capital and labor if $\psi < \epsilon$. One can see this by repeating the preceding analysis using ψ in place of φ (that is, using the long-run production function). This suggests that infrequent adjustment of the threshold task is important. With this in mind, we next sketch the model's dynamics when x is endogenous.

[37]The reason this is a local, rather than global, result is that there is also a scale effect. When the initial threshold task x is far from the optimum, a change in x induces not only capital-labor substitution but also an increase in TFP. However, the scale effect is second order when x is near x^*, while the substitution effect is first order. Thus, for an initial x near the optimum, the substitution effect always dominates.

[38]Recall from Proposition 3 that $\psi > 1$.

Proposition 4 describes the direct effect of an exogenous change in the allocation of factors to tasks on the inaction region. In a more realistic setting, the plant would endogenously choose x. In this case, what the Proposition suggests is that accounting for opposite-direction adjustment is much easier in a model with only *occasional* changes in x. For example, consider a case where $\varphi \approx 1$ and $1 < \psi < \epsilon$. Suppose that the plant experiences an (unanticipated) increase in capital-augmenting technology ξ_k. If x cannot change, the plant responds to the shock with an elasticity of substitution φ. We know from the discussion of capital-augmenting technology above that the parallelogram moves up and to the right, as in panel B of Figure 6. Figure 7 also illustrates this: the blue, dotted parallelogram lies to the northeast of its original, dashed counterpart. If the plant adjusts both factors, it will invest and hire. If x can adjust freely, on the other hand, the parallelogram also moves up and to the right; however, because $\psi > \varphi$, this movement is more rightwards and less upwards than before, so we see more investment and less hiring. In Figure 7, the solid parallelogram lies to the northeast of the dashed version, but to the southeast of the dotted one. Neither if x is fixed nor flexible would we see employment fall.

Now suppose that the firm only receives an opportunity to adjust x some time after the initial increase in ξ_k. Proposition 4 shows that, when the adjustment to x occurs, the firm "bridges the gap" and raises its capital intensity. The shift from the dotted parallelogram downwards and to the right to reach the solid parallelogram, then, delivers in a natural way an episode where the plant invests and reduces employment.

5 Dynamic model

We now embed the tasks-based model of production into a dynamic stochastic environment that incorporates frictions in adjusting capital, employment, and the threshold task, x. The friction on adjusting the task threshold generates the infrequent adjustments in x which we have just argued are important in generating opposite-direction factor adjustments.

We assume the task-based production function (9). This specification imposes that all tasks may be performed by either capital or labor, and, from Proposition 3, this implies that $\psi > 1$. As we shall see, the data will indeed 'ask' for a value of $\psi > 1$, at least if the factor-augmenting technology is capital-augmenting. (9) also allows for labor-augmenting technology, but our identification of ψ

in this case is much less sharp, as we discuss in Section 6. Thus, in what follows, we will simply proceed under the assumption $\psi > 1$. We defer a discussion of the stochastic process for technology until Section 6.

We continue to assume that the plant faces an isoelastic product demand schedule (2). To facilitate the analysis, it is helpful to impose that the demand shifter, ζ, follows a geometric random walk (as we assumed implicitly in the analysis of the canonical model):

$$\zeta' = \zeta e^{\varepsilon'}, \qquad \varepsilon' \sim N\left(-\frac{1}{2}\sigma_{\varepsilon}^2, \sigma_{\varepsilon}^2\right).$$

This is the only factor-neutral shock in the model. That it follows a random walk means we can normalize the firm's problem by ζ, and define factor demands in terms of k/ζ and n/ζ.

We allow for capital to depreciate at a positive rate δ per period. The wage w is taken as given; we discuss wages further in Section 7 below.

The factor adjustment costs with respect to capital and labor take the by-now familiar form (4). The choice of the threshold task x is also subject to a friction. There is little literature to guide how this friction is modeled. We have chosen a highly tractable approach that is, by now, very familiar in macroeconomics. In particular, we follow Calvo (1983): with probability λ the plant is constrained each period to set x at the same value x_{-1} used in the previous period, while with probability $1 - \lambda$ the plant can choose any $x \in [0, 1]$ at no cost. One way to rationalize this approach is to note that changes in capital intensity might require new machinery that is compatible with that plant's product mix and/or production process. The plant will adopt the machinery when it becomes available, but there is likely some randomness as to when that occurs.

We now combine these pieces in order to state the plant's optimization problem. The general problem involves six state variables: k_{-1} and n_{-1}, the plant's capital and labor inherited from last period; the factor-neutral and factor-augmenting technologies, ζ and $\xi \equiv (\xi_k, \xi_n)$ respectively; and the threshold task, x. Let ϱ be the discount factor applied to future profits. Then, if the plant does not have the opportunity to adjust x, its present value is

$$\begin{aligned}\Pi^0\left(k_{-1}, n_{-1}, x_{-1}, \xi, \zeta\right) = \max_{k,n} \Big\{ & \zeta^{1/\epsilon} F(k, n, x_{-1})^{\frac{\epsilon-1}{\epsilon}} - wn \\ & - \mathcal{C}_k\left(k, k_{-1}\right) - \mathcal{C}_n\left(n, n_{-1}\right) + \varrho\, \mathbb{E}_{\xi', \zeta'}\left[\Pi\left(k, n, x_{-1}, \xi', \zeta'\right)\right] \Big\},\end{aligned} \tag{16}$$

and if the plant does adjust x,

$$\begin{aligned}\Pi^{\Delta}\left(k_{-1}, n_{-1}, \xi, \zeta\right) = \max_{k,n,x} \Big\{ & \zeta^{1/\epsilon} F(k,n,x)^{\frac{\epsilon-1}{\epsilon}} - wn \\ & - \mathcal{C}_k\left(k, k_{-1}\right) - \mathcal{C}_n\left(n, n_{-1}\right) + \varrho\, \mathbb{E}_{\xi',\zeta'}\left[\Pi\left(k,n,x,\xi',\zeta'\right)\right] \Big\}. \end{aligned} \tag{17}$$

The function Π is simply a weighted average over Π^0 and Π^{Δ}, where the weight attached to Π^0 is the probability, λ, of not adjusting the threshold task next period:

$$\mathbb{E}_{\xi',\zeta'}\left[\Pi\left(k,n,x,\xi',\zeta'\right)\right] = \lambda \cdot \mathbb{E}_{\xi',\zeta'}\left[\Pi^0\left(k,n,x,\xi',\zeta'\right)\right] + (1-\lambda) \cdot \mathbb{E}_{\xi',\zeta'}\left[\Pi^{\Delta}\left(k,n,\xi',\zeta'\right)\right]. \tag{18}$$

The adjustment dynamics of the model are qualitatively similar to those of the simpler case of a one-time shock, discussed in Section 4.2. For example, an increase in ξ_k, for given x, shifts the parallelogram, with the magnitude and direction of the shift mediated by the substitution and scale effects. It also causes the plant's desired capital intensity to increase. When the opportunity to adjust x arrives, the parallelogram shifts to the southeast, as in Proposition 4, and this may cause the plant to increase capital and reduce employment. Web Appendix D.1 provides a more formal treatment of the policy rules.

6 Estimation

In this section, we explore whether the tasks-based framework of Section 5 is able to provide a satisfactory account of the factor adjustment dynamics we observe in Korean data. Since we wish to remain agnostic on how factor-augmenting technology should be modeled, we consider in turn the possibilities of capital- and labor-augmenting technical progress, in Sections 6.1 and 6.2 respectively. In each case, we attempt to identify structural parameters of the model that enable it to reproduce several salient moments of our data. We discuss our identification strategy, summarize our results, and note the limitations of our approach. Finally, we relate our results to the literature.

6.1 Capital-augmenting technology

Preliminaries. We simulate our model at an annual frequency.[39] We treat capital as the numeraire and normalize its purchase price to 1. A number of parameters are best informed by data outside of our panel. First, we set $\varrho = 0.95$ to be consistent with an annual real interest rate of 5.2 percent in Korea over our sample period, 1990-2006. Second, we set the annual depreciation rate to 10 percent, a standard value.[40] Third, we fix $\epsilon = 4$ as discussed in Section 4.2.

The law of motion for ξ_k is given by (19), which says that technology follows a geometric AR(1),

$$\ln \xi_k' = \rho_\xi \ln \xi_k + \varepsilon_\xi', \qquad \varepsilon_\xi' \sim N\left(0, \sigma_\xi^2\right). \tag{19}$$

In general, factor-augmenting technology includes an aggregate component. Our model, though, is designed to speak more narrowly to plant-level dynamics. In our view, then, a more suitable interpretation of ξ_k is as an index of a plant's capital productivity relative to the aggregate trend. That is why we specify factor-augmenting technology as a stationary process, with the idea that a plant may fall behind the mean at some times and exceed it at others.[41]

There are 9 parameters to estimate. One is the wage level, w.[42] Two parameters relate to the production function: φ and ψ, the short- and long-run elasticities of substitution, respectively. Three relate to the stochastic processes for product demand, ζ, and the capital-augmenting technology, ξ_k. In particular, we recover the standard deviations of the innovations to product demand, σ_ε, and capital-augmenting technical change, σ_ξ, in addition to the persistence of ξ_k, given by ρ_ξ. Finally, three parameters determine the extent of adjustment frictions on capital, labor, and the organization of production. These are the resale price of capital, p_s; the cost of labor adjustment, $c \equiv c^+ = c^-$; and the frequency of adjusting the task threshold, $1 - \lambda$. We impose symmetry of the employment adjustment costs since, to a good approximation, only the wedge $c^+ + c^-$ matters for the

[39]The results of Section 8.3 suggest that the results are likely to be robust to this choice.

[40]We can also allow for attrition of workers, modeled symmetrically to depreciation: in the absence of adjustment, employment takes the value $(1 - \delta_n)n_{-1}$, where δ_n is an exogenous constant attrition rate. However, in the benchmark model of tasks, we set $\delta_n = 0$. We consider robustness to this choice in Web Appendix F.4.

[41]It is possible to establish this normalization formally in a simpler setting including only capital and omitting factor adjustment costs. That is, if capital-augmenting technology contains an aggregate component Ξ, then the model can be normalized by this and re-expressed in terms of the plant's technology relative to Ξ (in which case Ξ will now merely scale the firm's profits).

[42]The numeraire is the price of investment goods; thus, $1/w$ is a measure of average labor productivity.

distribution of employment adjustment.[43] In summary, we have a list $\theta \equiv (w, \varphi, \psi, \sigma_\varepsilon, \sigma_\xi, \rho_\xi, p_s, c, \lambda)$ of 9 structural parameters.

We estimate θ via the method of simulated moments. We summarize this briefly, as it is now well known. Given a guess for θ, we solve the model by value function iteration. The optimal policy functions for k, n, and x are then used to simulate the behavior of a panel of 50,000 plants for 100 years (the moments are virtually invariant to the addition of more plants). We drop all but the last 20 years of data, so the number of years is similar to what we have in our Korean panel. We then calculate several moments using the simulated data. We repeat this with a new batch of sequences for ζ and ξ_k, and the moments are then averaged across the simulations. The parameter vector θ is updated as needed to minimize the (unweighted) norm of the distance between the simulated and empirical moments.[44]

Identification. Table 2 reports the moments we target and their values in our Korean dataset. In several cases, the relevance of a particular moment to a parameter is especially clear. For instance, the choice of the wage is informed by the observed labor share. Also, the frequencies of adjusting capital and employment are particularly informative about the magnitudes of the factor adjustment frictions, c and p_s. In addition, the variances of the individual distributions of factor adjustments—namely, the variances of i/k_{-1} and $\Delta \ln n$—reflect (modulo c and p_s) the variances of the innovations to product demand (σ_ε^2) and capital-augmenting technical change (σ_ξ^2).

Several moments relate more directly to the *joint* dynamics of capital and labor. For instance, we target the share of plant-year observations in which $i/k_{-1} > 0.1$ and $\Delta \ln n < 0$, as well as the average size of the associated decline in employment. The first of these moments speaks, in part, to the probability, $1 - \lambda$, that a plant is able to adjust its capital intensity. The second contributes further information about the sizes of shifts in both product demand and technical change.

Information regarding the cross-sectional dispersion in the capital-labor ratio helps identify ψ. To see why, suppose that a plant receives an opportunity to adjust x and decides, as a result, to re-optimize its factor demands. Given that x can be adjusted flexibly, the plant's production function exhibits an elasticity of substitution between k and n equal to ψ. In the absence of adjustment frictions on capital and labor, this would imply that the plant would exhibit an elasticity of $\ln k/n$

[43] Roys (2014) provides a useful discussion of this issue.

[44] We use the identity, rather than the optimal, weight matrix because the latter has inferior finite-sample properties (Altonji and Segal, 1996).

	Moment	Data	Model		
			k-augmenting	n-augmenting	
				Fixed x	Variable x
1.	Std. dev. of $\Delta \ln n$	0.220	0.227	0.245	0.241
2.	Std. dev. of i/k_{-1}	0.505	0.504	0.493	0.489
3.	Labor share	0.495	0.495	0.479	0.490
4.	Persistence of $\ln k/n$	0.638	0.637	0.634	0.621
5.	$\text{Var} \ln k/n$ ("spikes")	0.239	0.240	[0.155]	[0.185]
6.	$\text{Var}\, \Delta \ln k/n$ ("non-spikes")	0.011	0.011	[0.012]	[0.013]
7.	$\text{Prob}[\Delta n < 0 \,\vert\, i/k_{-1} > 0.1]$	0.360	0.359	0.330	0.346
8.	$\mathbb{E}[\Delta \ln n \,\vert\, i/k_{-1} > 0.1, \Delta n < 0]$	−0.186	−0.183	−0.169	−0.123
9.	$\mathbb{E}[\Delta \ln n \,\vert\, i/k_{-1} > 0.1, \Delta n > 0]$	0.208	0.202	0.202	0.240
10.	$\text{Prob}[\Delta n = 0]$	0.141	0.141	0.135	0.146
11.	$\text{Prob}[i = 0]$	0.478	0.479	0.472	0.469

Table 2. Moments in data and estimated models.

Notes: This table reports the values of the moments in the data and the values generated by the k- and n-augmenting models described in Sections 6.1 and 6.2 respectively. The moments in brackets in the columns for n-augmenting technology were not targeted in the estimation; see the text for discussion. Moment 4 is measured as the least squares coefficient in a regression of $\ln k/n$ on its own lag. Moment 5 is the variance of $\ln k/n$ among observations where investment, i, is positive; the change in employment, Δn, is negative; and there is a "spike" in $\Delta \ln k/n$ of at least 20 log points. Moment 6 is the variance of $\Delta \ln k/n$ among observations where i and Δn are same-signed; and the absolute change in $\ln k/n$ is no more than 20 log points. See text for further information on the other moments.

with respect to $\ln \xi_k$ equal to $\psi - 1$. While this result will not hold exactly when $c > 0$ and $p_s < 1$, still even in this case it is intuitive that, given the variance of ξ_k, the variance of $\ln k/n$ among x-adjusting plants helps identify ψ. But how do we isolate the set of plants that adjust x?

Proposition 4 suggests that plants that adjust their factors in opposite directions are more likely to have adjusted the threshold task. Therefore, we use the variance of $\ln k/n$ among plants where i/k_{-1} and $\Delta \ln n$ are (strictly) opposite-signed to help target ψ. Or, more specifically, we use the subset of these observations where $\Delta \ln k/n$ is "large." This is because plants which do not change x may still invest and reduce employment, either if the short-run elasticity φ is sufficiently far from unity (so that the inaction parallelogram can move southeast, as in panel A of Figure 6), even when x is fixed) or if a negative demand shock arrives at the same time as a positive shock to capital-augmenting technology. However, the changes in $\ln k/n$ in these episodes tend to be relatively small because scale and substitution effects offset. At our estimated parameters, we found that (conditional on i/k_{-1} and $\Delta \ln n$ having opposite signs), a sufficiently large change in the capital-labor ratio—specifically, if $|\Delta \ln k/n| > 0.2$ —was always associated with a change in x. We refer to these plant-year observations as "spikes" in k/n.

Next, to distinguish ψ from φ, again consider the static case. If a plant does not adjust its the threshold task, then its production technology exhibits an elasticity of substitution equal to φ. In the absence of adjustment frictions on capital and labor, its first-order conditions would imply that $\Delta \ln k/n = (\varphi - 1)\, \Delta \ln \xi_k$.[45] While this will again not hold exactly when $c > 0$ and $p_s < 1$, it still suggests that the variance of $\Delta \ln k/n$ among plants where x is *not* adjusted will be informative about φ. The model suggests how to select these plants in the data. For starters, plants where i/k_{-1} and $\Delta \ln n$ are same-signed are more likely to have not adjusted x. Second, there is good reason to focus on the sub-sample of these latter plants where $|\Delta \ln k/n|$ is "small." The reason is that employment can rise slightly even if x is increased if product demand increases sufficiently. But in these instances, capital will increase substantially relative to labor; $\Delta \ln k/n$ will be big. Therefore, to isolate cases where x did not change, we focus on observations where i/k_{-1} and $\Delta \ln n$ are same-signed and $\Delta \ln k/n$ is "small," in the sense that $|\Delta \ln k/n| < 0.2$ (the complement to the criterion above). We compute the variance of $\Delta \ln k/n$ among these, referred to as "non-spikes" in Table 3.

Other information from the joint dynamics also bears on φ and ψ. For instance, φ also partially governs the size of employment expansions when capital and labor move in the same direction. In addition, among investing and firing plants, a greater long-run elasticity of substitution amplifies the decline in employment. Hence, the average value of $\Delta \ln n$ in these episodes helps pinpoint ψ.

The only parameter that remains is ρ_ξ, the persistence of the capital-augmenting technology. Factor-augmenting technical change is a key driver of the capital-labor ratio. Therefore, ρ_ξ bears strongly on the persistence of $\ln k/n$. This is why we include as a moment the projection of $\ln k/n$ on its own lag.[46]

Results. The second column of Table 2 reports the model-generated moments. Overall, the fit of the model is very encouraging. Although the model is slightly over-identified, it replicates all of

[45] The first difference eliminates $x = x_{-1}$ from the expression.

[46] Plants have permanently different capital-labor ratios for reasons absent in the model. To control for this, we calculate the mean, $\mathbb{E}[\ln k_i/n_i]$, of the capital-labor ratio for each plant i and then project $(\ln k_i/n_i - \mathbb{E}[\ln k_i/n_i])^2$ on a constant. The intercept yields an estimate of the variance of $\ln k/n$, adjusting for permanent heterogeneity. We also include plant fixed effects in the AR(1) regression when estimating persistence.

Parameter	Meaning	k-augmenting	n-augmenting	
			Fixed x	Variable x
φ	Short-run elasticity	1.402	0.01	0.025
		(0.009)	-	-
ψ	Long-run elasticity	5.600	n.a.	1.500
		(0.049)	-	-
σ_ϵ	Std. dev. of neutral innovations	0.238	0.158	0.121
		(0.001)	(0.005)	(0.008)
σ_ξ	Std. dev. of factor-augmenting tech.	0.284	0.458	0.463
		(0.003)	(0.002)	(0.002)
ρ_ξ	Persistence of factor-augmenting tech.	0.292	0.651	0.696
		(0.007)	(0.004)	(0.003)
λ	Probability of $\Delta x = 0$	0.178	1.000	0.933
		(0.004)	-	-
w	Real wage	0.126	0.247	0.079
		(0.001)	(0.006)	(0.001)
p_s	Resale price of capital	0.932	0.766	0.760
	(% of purchase price)	(0.001)	(0.001)	(0.003)
c	Labor adjustment cost	0.014	0.092	0.046
	(% of annual wage)	(0.0001)	(0.002)	(0.001)

Table 3. Parameter estimates.

Notes: This table presents estimates of the model's structural parameters for the k- and n-augmenting models in Sections 6.1 and 6.2 respectively. Standard errors are in parentheses. Dashes (-) indicate parameters for which standard errors were not computed; either because the parameter is fixed a priori or because (in the case of φ in the n-augmenting fixed-x specification) the estimated value is the lower bound of the range considered.

the moments quite precisely. We call attention, in particular, to the model's implications for the joint dynamics. It reproduces the frequency of years when investment is positive and employment falls, and yields a decline in employment in these episodes that is similar to what is observed. Moments of the marginal distributions of i/k_{-1} and $\Delta \ln n$ are also reproduced accurately: the model captures the mass at zero, and the dispersion around zero, in each case.

Table 3 reports the estimated parameters. We discuss first the parameters most directly related to the production technology. The short-run elasticity, φ, is estimated to be 1.4. Even in the short run, then, factors are gross substitutes (but $q-$complements). Still, φ is substantially smaller than the long-run elasticity, ψ, which is found to be 5.6. Moreover, changes to plants' capital intensities, via changes in x, happen reasonably often, as λ (the probability of not adjusting) is slightly less than 0.2.

The variability of the threshold task, x, and the distinction between φ and ψ are key to the model's ability to engage the joint dynamics. To see this, we have re-solved the model with $\lambda = 1$ (without changing the other parameters). This implies that plants may not alter their capital intensities (the probability of adjusting x is zero), and so the long-run elasticity plays no role. In this case, the model generates hardly any periods of investment and negative employment growth. Indeed, the probability of $\Delta \ln n < 0$ given $i/k_{-1} > 0.1$ is about 1 percent.

A few other parameter estimates deserve special mention. First, our estimates of the factor adjustment costs have precedent in the literature. For instance, the cost of adjusting employment, expressed in Table 3 as a share of the annual wage, is similar to, though slightly smaller than, that in Bloom (2009). Estimates of the capital adjustment friction are more varied in the literature. Ours is much smaller than in Bloom (2009), who finds a discount of $1 - p_s = 0.34$, but comparable to Cooper and Haltiwanger's (2006) estimate of $1 - p_s = 0.025$. Second, the value of $\rho_\xi = 0.29$ is somewhat low. The reason why ρ_ξ is driven so low is that a persistent process "blows up" the variance of investment rates and $\ln k/n$. One could offset this with a higher λ, so that capital intensity (x) is less volatile. But this would make it difficult to replicate the joint dynamics of factor adjustment. The estimator instead lowers ρ_ξ and relies, in part, on factor adjustment frictions to induce persistence in the capital-labor ratio. Recall that ξ_k should be interpreted as the deviation of the plant's capital-augmenting technology from the mean.

Finally, to sharpen our intuition about how the parameters are identified, we have also analyzed the effects of perturbations to the parameters around their estimated values. Numerical results are reported in Web Appendix D.2. We summarize them here. First, as we anticipated, reductions in φ imply a more disperse distribution of $\Delta \ln k/n$ among plants with same-signed changes in capital and labor. Second, the dispersion in $\ln k/n$ among plants with *opposite*-signed changes in capital and labor increases with ψ, also as anticipated. Lastly, changes in λ have significant, and intuitive, effects. Increases in λ greatly reduce the likelihood of opposite-signed changes in factor demands. In addition, since changes in the threshold drive big changes in capital, fewer changes in x mean a much smaller variance of i/k_{-1}.

6.2 Labor-augmenting technology

Thus far, we have focused on capital-augmenting technology. Labor-augmenting technology is arguably the more common choice. Here, we take up two issues. First, we discuss why the specification of the model with labor-augmenting technology is, in general, difficult to identify. Second, we discuss how, nonetheless, the general tasks-based framework nevertheless adds value as a descriptive model of plant dynamics when there is labor-augmenting technology.

First, identification under labor-augmenting technology is more challenging than under capital-augmenting technology. To see why, suppose we fixed the threshold task x and tried to account for our data. Consider two moments in particular: the variance of the investment rate (i/k_{-1}) and the co-occurrence of $\Delta \ln n < 0$ and $i/k_{-1} > 0$. Section 4.2 illustrated why, under either n- or k-augmenting technology, the second of these moments may be replicated with a (very) low value of φ even if $\lambda = 1$, that is, even if the threshold is fixed. However, under k-augmenting technology, it turns out that under the values of φ for which this occurs, the volatility of investment, relative to the volatility of employment growth, will be either much too low (for φ near 0) or much too high (for φ above ϵ). For example, consider the case where φ is near 0, that is, the production function is nearly Leontief. Intuitively, after an increase in ξ_k, a plant wants to bring n up and into line with $\xi_k k$. This implies a relatively large increase in n, but a more muted change in k. To illustrate, consider the case with no adjustment frictions for simplicity. The elasticities with which k and n respond are respectively given by $\varphi - 1 + (\epsilon - \varphi)s$ and $(\epsilon - \varphi)s$; for $\epsilon = 4$, $s = 1/2$, and $\varphi \to 0$, this implies that capital adjustment is half as volatile as labor adjustment.[47] In our data, on the other hand, the standard deviation of investment (0.505) is more than twice that of employment growth (0.220).

This tension between the marginal and joint distributions of factor adjustments is resolved under k-augmenting technology by variation in the threshold: a combination of a higher φ and a choice of $\lambda < 1$ enables the model to generate both a large dispersion of investment rates and a co-occurrence of $\Delta \ln n < 0$ and $i/k_{-1} > 0$. In particular, a relatively high φ is called for. As we saw in Section 4.2, a φ of this magnitude does not generate any episodes in which employment is reduced as investment

[47] However, inference from the frictionless case is only partial: in particular, in the absence of adjustment frictions, the firm would never adjust capital and labor in opposite directions for these parameters. To contrast, under frictions, there will be plants whose initial n is larger than "desired" and thus a fall in ξ_k will lead to both investment and separations.

is undertaken. Thus, the "need" to replicate this latter moment of the data helps identify both ψ and λ.

Under labor-augmenting technology, however, this tension between the marginal and joint distributions is much less severe, and thus identification is much less sharp. In this case, a low φ helps induce episodes of $\Delta \ln n < 0$ and $i/k_{-1} > 0$. In addition, and unlike k-augmenting technology, labor-augmenting technology *amplifies* the variance of capital adjustments relative to employment adjustments.[48] Thus, it is much harder to identify a role for threshold (x) adjustments given the moments we are able to use from our plant-level data.

To illustrate this, we have fixed the threshold task x by imposing that $\lambda = 1$, solved the model under n-augmenting technology, and looked for a parameter vector that best fits our plant-level data. For this exercise, ξ_n follows a geometric AR(1) process analogous to (19). We then select six parameters $(w, \sigma_\varepsilon, \sigma_\xi, \rho_\xi, p_s, c)$ to target nine of the eleven moments listed in Table 2.[49] The resulting moments and parameters are reported in the second columns in Tables 2 and 3. This standard CES model fits quite well. In particular, and as foreshadowed in Section 4.2, it yields a realistic co-occurrence of $\Delta \log n < 0$ and $i/k_{-1} > 0.1$ if $\varphi \to 0$.[50] This shows that, if φ is free to fall to the Leontief limit, there would be relatively little variation "left over" in the moments of our data to identify a distinct role for the threshold, x.

Some of the parameter estimates from this specification do deserve some comment. First, it is noticeable that the neutral shock plays a lesser role in this specification than in the k-augmenting case. This is because because, as discussed above, the n-augmenting shock alone generates realistic capital and employment adjustments, with less need for the (same-direction) adjustment episodes generated by movements in ζ. Second, the adjustment costs are larger than in the case of k-augmenting technology, but this is a natural consequence of the much higher standard deviation and autocorrelation of the shocks to the relevant factor-augmenting technology: the frictions damp down the factor adjustments implied by the larger shocks.

[48] In the absence of adjustment costs on capital and labor, the elasticity with which capital responds to an increase in ξ_n is $(\epsilon - \varphi)(1 - s)$, while that for labor is $(\varphi - 1) + (\epsilon - \varphi)(1 - s)$. For $\epsilon = 4$, $s = 1/2$, and $\varphi \to 0$, now capital responds twice as elastically.

[49] We do not target the moments that relate to the variances of $\ln k/n$ and $\Delta \ln k/n$. These moments were introduced to help identify parameters ψ and λ that are unique to the general, flexible-x model of tasks.

[50] The estimated value of φ is 0.01, which is the lower bound imposed in the estimation process. We show in Web Appendix E that increasing φ even a little quickly reduces the frequency of episodes in which the plant invests and reduces employment, and also rapidly shrinks the sizes of these employment reductions.

What happens if we relax our assumption that the threshold x is fixed, and allow for flexibility in x? As discussed, given our current data, it is very difficult to disentangle the roles of φ, ψ, and λ if we do so. Therefore, we instead propose a simple exercise that tries to explore the potential of such an integrated model. The idea is to fix ψ and λ and then identify a parameter vector that best fits the data. For the purposes of illustration, we fix $\psi = 1.5$, which is at the high end of the range of estimates in Karabarbounis and Neiman (2014). We also fix $\lambda = 0.933$, so that one-half of firms adjust their threshold within a decade. With these parameters in place, we then adjust the remainder of the structural parameters to target the same nine moments targeted in the fixed-x exercise described above.

The resulting moments and parameters are reported in the rightmost columns of Tables 2 and 3 respectively. The results show promise but also reveal a challenge to this integrated framework. The fit is good with respect to most moments, for instance, the co-occurrence of $\Delta \ln n < 0$ and $i/k_{-1} > 0.1$. However, the average size of the decline in n in the latter episodes is too small. There appears to be a good reason for this. If ξ_n increases, the low φ means that, absent labor adjustment frictions, the firm wants to reduce employment. However, the firm now foresees that it will, with some probability, have the opportunity to reallocate tasks toward the now-abundant factor (labor). Since it is costly to reverse its employment decision, the firm attenuates the initial decline. Thus, even though $\lambda < 1$ introduces additional substitution of capital for labor, it also implies smaller declines in n in periods where x is fixed and investment is undertaken.

6.3 Discussion

Let us take stock. We have presented three several specifications of our model—with capital-augmenting technology shocks, and with labor-augmenting shocks both with fixed and variable task allocation—each of which can account for (at least most of) the moments reported in Table 2. How do we assess these three specifications?

First, it is remarkable in the case of k-augmenting technology that our tasks-based model is so well able to account for the data. No pure (fixed-x) CES model with k-augmenting technology is able to account for the frequency and magnitude of opposite-direction factor-adjustment, much less jointly with the additional nine moments we target in Table 2; the flexible-x model discussed in Section 6.1 above, on the other hand, does an excellent job.

In the case of n-augmenting technology, on the other hand, the advantage of the flexible-x model over its fixed-x counterpart is not so much that it fits the data significantly better—as noted, identification in this case is challenging—but because of its very different implication for the elasticity of substitution. In fact, the most remarkable difference between the three variants of our model lies in their implications for this key elasticity. For example, the model with k-augmenting technology requires a very *high long-run* elasticity ψ (as well as a large difference between the short- and long-run elasticities) to account for episodes in which capital and labor are adjusted in opposite directions. It interprets these episodes as occasions where the threshold task x changes. The fixed-x model with n-augmenting technology, on the other hand, generates these episodes using a very *low short-run* elasticity φ, such that the substitution effect dominates the scale effect. A good way to assess the various specifications of our model, then, is to look at evidence from outside our model on the elasticity of substitution.

In empirical work, one common way to summarize the effect of factor price changes uses the factor expenditure ratio, rk/wn, where r is the implicit rental rate on capital (ignoring irreversibilities) and wn is the wage bill. In a standard CES model, the elasticity of rk/wn with respect to the factor price ratio w/r is the elasticity of substitution, minus 1. Using cross-metropolitan variation in wages in U.S. manufacturing and assuming such a model, Raval (2014) and Oberfield and Raval (2014) estimate plant-level elasticities of substitution of about 0.5, that is, an elasticity of the factor expenditure ratio of around -0.5. Chirinko, Fazzari, and Meyer (2011) analyze a cross-section of Compustat firms and uncover a factor expenditure ratio elasticity of about -0.6.[51]

We can assess different variants of the tasks framework in light of this evidence on the factor expenditure ratio. It turns out that the elasticity of rk/wn implied by our simulations for the response T years after a small permanent factor price change is well approximated by $\lambda^T(\varphi - 1) + \left(1-\lambda^T\right)(\psi - 1)$, which is intuitive: after T years, a fraction λ^T of firms have not had the opportunity to adjust x, so react to the factor price change with substitution elasticity φ, while the remaining firms have adjusted x and so respond with elasticity ψ. If we consider the response at a 10-year horizon, as do Raval (2014) and Oberfield and Raval (2014), this suggests our k-augmenting specification implies an elasticity of rk/wn close to $\psi - 1 = 4.600$ (since $\lambda^{10} = 0.178^{10} \approx 3 \times 10^{-8}$ is

[51] Oberfield and Raval (2014) allow for reallocation of factors across heterogeneous plants in response to relative price changes, and find a higher aggregate elasticity of substitution around 0.7. Chirinko (2008) surveys research on the elasticity of substitution. He highlights the challenges to estimating the long-run elasticity in particular.

very small). Of course, the n-augmenting specification with fixed x implies a corresponding elasticity of $\varphi - 1 = -0.99$. Each of these is remarkably far from standard estimates. For the variable-x specification under labor-augmenting technology, on the other hand, the corresponding elasticity should be around $0.933^{10} \cdot (0.025 - 1) + \left(1 - 0.933^{10}\right) \cdot (1.500 - 1) = -0.24$; in fact, we find a value of -0.26 in our simulations.

In summary, the advantage of allowing for a flexible task allocation under n-augmenting technology is that it allows our model to match at least in general terms the striking *short-run* factor adjustment we see in our data, while at the same time being consistent with estimates from other data sets of the *long-run* elasticity of substitution. We emphasize that reconciling these two sets of facts is crucially reliant on the most novel feature of the tasks framework, which is to disentangle the short- and long-run substitution elasticities in a tractable way.

We conclude with a few thoughts on the variant with k-augmenting technology. As noted above, this model requires a long-run elasticity above 5 to fit the data. However, we still regard its study as highly informative. First, this version of our model in fact does the best job of any of the three variants we have discussed of matching our moments. Moreover, there is certainly not universal agreement that the long-run elasticity of substitution is 1 or below. Some authors have argued that medium-run changes in factor shares in many advanced economies are highly suggestive of a relatively high long-run elasticity of substitution, though few would likely defend an elasticity in the neighborhood of 5.[52] Histories of long-run growth have also emphasized how more advanced machinery has often replaced labor in tasks (Habakkuk, 1962; Acemoglu, 2010). These theories may still generate stable aggregate factor shares in the long run but via a process of structural change, whereby labor migrates from high- to low-ψ industries (Zeira, 2006).

7 Worker heterogeneity

The baseline tasks-based model described in Section 5 and estimated in Section 6 assumed that the wage paid by the plant is fixed. Yet in fact the average wage paid in our data *in*creases during the episodes we find most striking, namely, when $i/k_{-1} > 0.1$ and $\Delta \ln n < 0$. The average real wage rises by around 11 percent when plants invest and reduce employment, or around 6.5

[52] Caballero and Hammour (1998) do indeed entertain such a high elasticity. Karabarbounis and Neiman (2014) find evidence of a more modest elasticity of around 1.25.

percentage point in excess of the sample mean. This is important in accounting for why labor's share at these plants falls only by 3 percentage points, despite the large decline in labor demand. In this section we show that an extension of our baseline model which allows for worker heterogeneity can be consistent with these facts. Building on the ideas of Griliches (1969) and Goldin and Katz (1998), we interpret the behavior of the wage in episodes of investment and employment reduction as a shift in the composition of the workforce. When the plant invests, it sharply decreases its employment of low wage workers, and partially offsets this by hiring higher-paid workers who are more complementary to the new capital.

Before we describe our extended model, an initial comment is in order. Why do we not try to account for plant-level wage movements via bargaining? After all, if the plant bargains with its workers over the division of the rents arising from the costs of adjusting employment, then shocks to demand or technology can affect wages. However, the most commonly used wage bargaining protocols in environments like this (for example, Stole and Zwiebel, 1996) have the property that the wage increases with the marginal product of labor. If the plant reduces employment, this suggests that labor's marginal product has fallen. It is therefore unclear whether standard wage bargaining models can deliver a useful explanation for wage increases when the plant invests and sheds workers.[53]

We incorporate worker heterogeneity into our model in a straightforward way. We assume that, when the firm assigns capital to a task, it must also assign some workers to that task. We refer to these workers as *technicians*. Capital and technicians must be combined in fixed proportions, as in Goldin and Katz (1998). This Leontief structure preserves the model's tractability.[54] If the firm uses k_i machines in task i, we assume that it must also assign $l_i = k_i/\eta$ technicians, where η is the number of machines used by each technician and is constant across tasks. The tasks to which capital and technicians are not assigned are performed by *assemblers*. We denote their quantity on task i by m_i. We assume that the productivity of the factors in the different tasks is still as in the

[53] Of course, other wage bargains are possible. For example, the bargain might transfer a share of *total* output to workers. The idea is that, if workers anticipate that machines may replace them, they negotiate for a share in the firm. But this implies very similar results as the approach we take, as long as the worker's bargaining power (the share of output she captures) is chosen to replicate the observed changes in average wages. We thank Andy Glover, whose comments inspired this exercise.

[54] See Krusell et al. (2000) for a general CES model in which one cohort of workers (the skilled) is a gross complement to capital and another cohort (the unskilled) is a gross substitute. We depart from Krusell et al. in two ways. First, we include a margin for adjusting capital intensity (x), which implies a clear distinction between short- and long-run elasticities. Second, we borrow from Goldin and Katz's Leontief structure since it delivers a more tractable model.

baseline model: that is, (6) continues to hold, with the new interpretation that a task performed by capital must now also be assigned the appropriate number of technicians.

Given this structure, we assume that the plant (as a whole) always uses η machines for each technician. Thus, if l is total employment of technicians, then $k = \eta l$.[55] In this case, the plant's factor inputs can be summarized simply in terms of its usage of machines, k, and assemblers, m. If we assume that the productivities of capital (and technicians) and of assemblers in the different tasks are as in part 1 of Proposition 3, then the plant's short-run production function is still given by (9), appropriately adjusted, that is,

$$F(k,m,x) = \left[x^{\Gamma}\left(\xi_k k\right)^{\frac{\varphi-1}{\varphi}} + (1-x)^{\Gamma}\left(\xi_m m\right)^{\frac{\varphi-1}{\varphi}}\right]^{\frac{\varphi}{\varphi-1}},$$

where Γ is as in Proposition 3. The remaining assumptions of the model hew closely to those of the baseline, and are described in more detail in Web Appendix D.3.1. We assume in particular that the adjustment costs which apply to technicians are equal to those which apply to assemblers. Also, for illustration, we work only with the case of capital-augmenting technology shocks, as in the model of Section 6.1, and we assume that $\psi > 1$ from the outset.

The augmented model has two extra parameters in addition to those used in the baseline model: w_l/w_m, the wage premium of technicians relative to assemblers, and η, the number of machines per technician. We fix the wage premium outside of the model. We identify occupations in Korean manufacturing that best fit our vision of technicians as carrying out the design, repair, and setting of machines as well as the development of manufacturing processes. The average monthly wage in this group is 54 percent higher than the wage among the remaining occupations in manufacturing, the most common of which is "labourer."[56] Accordingly, we set $w_l/w_m = 1.54$. (Web Appendix D.3.4 reports results for alternative parameterizations.) Given this value, we use w_m to target the labor share, just as we did with w in the baseline model.

After we have determined the wage premium w_l/w_m, the identification of η is intuitive. In particular, if η is low, many technicians must be hired when the plant becomes more capital intensive.

[55]In the absence of adjustment frictions, this property would follow immediately from cost minimization, but in general, it does amount to a restriction on the plant's strategies. For example, we rule out that the plant might respond to a bad demand or technology shock by "hoarding" technicians more than machines, in order to economize on adjustment costs.

[56]We use the Occupational Wages Around the World database (Oostendorp, 2012). Web Appendix D.3.2 discusses our choice of technician occupations in more detail.

Moment	Data	Model	
		Baseline	Heterogeneity
Std. dev. of $\Delta \ln n$	0.220	0.227	0.225
Std. dev. of i/k_{-1}	0.505	0.504	0.509
Labor share	0.495	0.495	0.485
Persistence of $\ln k/n$	0.638	0.637	0.634
$\mathrm{Var} \ln k/n$ ("spikes")	0.239	0.240	0.248
$\mathrm{Var}\, \Delta \ln k/n$ ("non-spikes")	0.011	0.011	0.011
$\mathrm{Prob}[\Delta n < 0 \mid i/k_{-1} > 0.1]$	0.360	0.359	0.342
$\mathbb{E}[\Delta \ln n \mid i/k_{-1} > 0.1, \Delta n < 0]$	−0.186	−0.183	−0.190
$\mathbb{E}[\Delta \ln n \mid i/k_{-1} > 0.1, \Delta n > 0]$	0.208	0.202	0.182
$\mathrm{Prob}[\Delta n = 0]$	0.141	0.141	0.112
$\mathrm{Prob}[i = 0]$	0.478	0.479	0.496
$\mathbb{E}[\Delta \ln w \mid i/k_{-1} > 0, \Delta n < 0]$	0.065	n.a.	0.067

Table 4. Moments for model with worker heterogeneity.

Notes: This table reports the values of the moments in the data and the values generated by the model with worker heterogeneity described in Section 7. For comparison, the results for the baseline k-augmenting specification are repeated from Table 2.

This means an increase in x is accompanied by a large increase in high-wage workers. As a result, the mean wage at the plant, $(w_l l + w_m m)/(l + m)$, will rise. We therefore use η to target the change in the mean wage in years when the plant invests and reduces employment. Note that, since we abstract from aggregate growth in the model, we choose to target the real wage change in excess of the sample mean, that is, 6.5 percent. The intuition behind the identification of the remainder of the parameters has already been discussed.

The results are reported in the right hand column of Table 4. Results for the k-augmenting baseline specification, as well as the target moments from the data, are repeated from Table 2 for convenience. The model's fit suffers slightly relative to the model with homogeneous labor with regards to the original eleven targeted moments. In particular, labor adjustments are a little too frequent and, given $\Delta \ln n > 0$, too small. However, the model largely accomplishes the objective we set: it reproduces several salient moments of the joint dynamics and generates realistic changes in the average wage.

This does require a few noteworthy changes to the parameters. We report the full list of parameters in Web Appendix D.3.3; here we comment on the most salient. The required estimated value of ψ increases from 5.600 to 7.063. This larger elasticity is required to generate the observed declines in $\Delta \ln n$ at plants which invest and fire, since the substitution for assemblers must be large enough to offset the coincident increases in technician demand. There are also larger k-augmenting innovations—σ_ξ increases from 0.284 to 0.359—which serves to generate the needed variation in x, and thus target the joint dynamics.

Lastly, how do we interpret η? Our estimate of the technicians per unit capital is 0.209, which implies that 20 percent of the workforce consists of technicians. This is less than the mean non-production share over our sample, which is about 30 percent. But as discussed below, it appears that even some non-production workers can be substituted by new computers and other capital. Thus, it seems plausible to us that the share of the workforce that is (strongly) complementary to capital, even in the long-run, might be smaller.

Given that we account for plant-level wage movements through compositional changes, it would be desirable to observe such composition changes directly. Unfortunately, data limitations preclude this. The only arguably relevant information we observe in our data is the division of plants' employment into production and non-production workers. Non-production status is sometimes treated as a proxy for skill.[57] However, the behavior of the two types of employment is rather similar. The share of plant-year observations in our data that involve both investment and declines in *non*-production employment is still rather high at 30.6 percent, only 5.4 percentage points less than for total employment (see Table 1). Moreover, the average decline in non-production employment in these periods is not any smaller than the decline among the general workforce reported in Table 1.[58] This does not rule out accounting for wage changes using compositional shifts; it merely requires that production or non-production status not be too tightly related to a worker's complementarity with capital. For instance, both less-skilled non-production workers (for example, in certain clerical and administrative jobs), as well as less-skilled production workers, can be replaced by more advanced computers and new machinery (Acemoglu and Autor, 2011).

[57]See, for example, Berman, Bound, and Machin (1998) and Berman and Machin (2000).

[58]There is an upward trend in the non-production worker share in our data. Since our model environment is stationary, we abstract from this.

Moment	Data	Model			
		Canonical (1)	Delivery lag (2)	Disruption (3)	Replacement (4)
$\text{Prob}[\Delta n < 0 \mid i/k_{-1} > 0.1]$	0.360	0.046	0.125	0.030	0.188
$\mathbb{E}[\Delta \ln(n) \mid i/k_{-1} > 0.1, \Delta n < 0]$	−0.186	−0.041	−0.035	−0.036	−0.111
Probability per quarter of delivery of new capital, $\varkappa$	0.33	n.a.	0.33	n.a.	n.a.
Share of revenue lost due to disruption, τ	n.a.	n.a.	n.a.	0.025	n.a.
Replacement share of total investment	0.6	n.a.	n.a.	n.a.	0.6

Table 5. Alternative hypotheses.
Notes: See main text for further details.

8 Alternative hypotheses

The canonical framework of Section 2 abstracted from some features of the economic environment that could potentially be important in accounting for the joint factor dynamics we find in Korean data. In this section, we study several important possibilities. Our goal is to investigate whether modifying assumptions made in the canonical framework that do *not* pertain to the production technology allows that model to deliver a full quantitative account of the factor adjustment distribution. Accordingly, we maintain the assumption of a Cobb-Douglas revenue function, as in (3), throughout.

We consider six candidate factors: wage bargaining, sticky product prices, time aggregation, lags in the delivery of machinery, disruption costs of investment, and required replacement investment. The last four possibilities seem especially promising, so we assess them quantitatively. Our objective is *not* to estimate the canonical model thus augmented against the full array of moments. Rather, we assess how important each mechanism is for replicating salient moments of the joint distribution when disciplined by external evidence on its importance. The results reveal whether the observed joint dynamics can emerge as mere "by-products" of more standard theories. For the last four possible factors, Table 5 reports the key moments that summarize the joint factor dynamics. As noted, these were not targeted in the calibration. Table F.1 in Web Appendix F lists the structural parameters we use for each exercise, and reports the targeted moments.[59]

[59]Web Appendix F.4 also reports results of additional sensitivity analysis that relaxes a few restrictions imposed in canonical model to retain tractability—for instance, the assumption that ζ follows a random walk.

8.1 Wage bargaining

The canonical framework assumed that the wage is fixed. However, given the costs of adjusting labor, there are rents associated with ongoing employment relationships, and it is natural to allow for the possibility that these might be shared between the plant and its workers by bargaining. Because there are decreasing returns to labor within the plant, the most natural candidate bargaining protocol is that of Stole and Zwiebel (1996) and Smith (1999). For this bargaining protocol, the wage takes a form that preserves the supermodularity of the plant's revenue function. Thus, allowing for this type of bargaining would not affect the predictions of the canonical model for factor dynamics. Moreover, even if there is a stochastic exogenous component to the wage, our conclusions are unaltered.[60] The reason is simple: if factors are complements, a higher wage triggers an increase in the capital-labor ratio at an adjusting firm, but reduces demand for each factor in *absolute* terms. Thus, factors still move in the same direction.[61] Web Appendix F.1 gives a more detailed discussion of wage bargaining.

8.2 Sticky product prices

We implicitly treated the product price as flexible. If the plant's product price is sticky, factor-neutral technical change can be contractionary (Basu, Fernald, and Kimball, 2006). The reason is that, if the plant does not lower its price in order to sell additional output, it does not need its current factors given the higher level of productivity. Thus, its factor demand declines. However, technology is contractionary in this case with respect to *both* capital and labor. It is hard to see, then, how this mechanism explains the extent of opposite-direction adjustment of the factors.

8.3 Time aggregation

The plant may decide its factor demand more often than the frequency of our data. This means that it could invest and hire at one instant, but then reduce employment significantly before the year's end. As a result, the net change in employment is negative. By this view, then, the observed data on joint factor adjustment are the result of *time aggregation*.

[60] In a bargaining model, this component would correspond to the worker's outside option.

[61] This is easy to verify analytically in a static model if the revenue function is Cobb-Douglas. We have found numerically that it remains true in the dynamic case.

To address this, we simulate the model at a quarterly frequency and then aggregate the simulated data to an annual frequency and compare to the Korean data. The discount and depreciation rates are adjusted appropriately. We also re-calibrate several parameters to target the factor adjustment distributions. The adjustment costs, c and p_s, are chosen to match the frequencies of adjusting capital and labor, and the standard deviation of the innovation to ζ, σ_ε, is chosen to replicate the average expansion in employment conditional on $i/k_{-1} > 0.1$ and $\Delta \ln n > 0$. We then ask if the model is able to reproduce the average *decline* in employment if $i/k_{-1} > 0.1$ and $\Delta \ln n < 0$.

As seen in Table 5, the canonical model reproduces neither the frequency of $i/k_{-1} > 0.1$ and $\Delta \ln n < 0$ episodes nor the average size of the decline in n in these periods. This reveals that time aggregation is relatively unimportant. This is so because, if a plant invests at least 10 percent of its capital, the associated increase in employment is typically large. Given the sizes of the shocks ζ and adjustment frictions (c, p_s), it is very unlikely that this increase is reversed within the same year.[62]

Before we proceed, we note that all of the exercises below build in time aggregation, that is, they are solved at a quarterly frequency and the simulated data are aggregated to an annual frequency. Thus, the effect of each mechanism considered below can be isolated by comparison with the results in column (1) of Table 5.

8.4 Delivery lags

The canonical model assumes that new machinery is installed immediately. Delivery and installation may take time, though. If the plant reports investment in the year in which the order is placed, then this delay does not disrupt the implications of the canonical model: orders, and thus recorded investment, should still perfectly predict positive employment growth. However, if plants record investment when the machinery is *delivered*, we could observe a decline in employment coincident with investment merely because productivity has fallen by the time the capital goods arrive.

To investigate this, we introduce delivery lags into the canonical model. In general, this can introduce several new state variables—one to track unfilled orders placed last quarter, another to track unfilled orders placed two quarters ago, and so on. To circumvent this, we assume that if

[62] If we condition instead on a net investment rate of 1 percent, then 18 percent of observations in the simulated data involve positive investment and a year-over-year net decline in employment. However, the average decline in employment among these plants is just 4 percent, which is less than one-fourth of the average in Korean data.

a plant orders new capital, the order is added to a backlog of unfilled orders. Each period, with probability $\varkappa$, the plant's *entire* order backlog is filled. We set $\varkappa = 1/3$, so the expected delivery lag is three quarters. This corresponds to the high end of the range estimated for the U.S. by Abel and Blanchard (1988). The model is otherwise identical to the canonical model, and the parameters c, p_s, and σ_ε are chosen in the same manner described just above.[63]

Results are reported in column (2) of Table 5. The share of plant-year observations that involve opposite movements in investment and employment is now 12.5 percent, far higher than in the canonical model but still just one-third of the empirical estimate. However, the mean decline in log employment, conditional on positive investment, is 3.5 points, which is less than one-fourth what we see in the data. This suggests that it is difficult to generate large reductions in employment purely as a by-product of delivery lags.

8.5 Disruption costs

Cooper and Haltiwanger (1993) observed that the installation of machinery may disrupt production. For instance, the assembly line may have to suspend work or slow its pace while new equipment is put in. This diminishes labor demand. Hence, *disruption* may generate periods of positive investment and net separations.

To capture this consideration, we assume that if the plant undertakes investment, it now faces a cost of adjustment equal to a fraction $\tau \in (0, 1)$ of its revenue. Thus, conditional on investment, the plant earns revenue $(1 - \tau) \cdot F(k, n, \zeta)$.[64] In periods of installation, the new capital is not yet operational, but the disruption it causes degrades the marginal product of labor.[65] However, if employment is costly to adjust, plants may nonetheless "hoard" labor, since workers' marginal product will be high once the new capital becomes active. Whether disruption costs can account for the co-occurrence of employment declines and investment is therefore a quantitative question.

We must parameterize τ. It takes only a small value of τ to induce a lot of inaction with respect

[63] The model of delivery lags is described more formally in Web Appendix F.2.

[64] Cooper and Haltiwanger (1993) did not allow for stochastic productivity, but a specification like this has been used often in the more recent dynamic factor demand literature. See, among others Cooper, Haltiwanger, and Willis (2004), Cooper and Haltiwanger (2006), and Bloom (2009).

[65] This assumption deviates from the canonical model, where new capital becomes productive immediately. We make this modification because a one-period delay gives the disruption cost the best chance of generating employment declines at times of investment.

to capital.[66] Yet if τ is set too low, installation will not disrupt the marginal product of labor. As a compromise, we set $\tau = 0.025$, which implies that investment is zero 70 percent of the time—an inaction rate 1.5 times that in the data.[67] Since τ now induces inaction, we can abstract from other forms of irreversibilities, and set $p_s = 1$.

Column (3) of Table 5 displays the results. The model with disruption generates no co-occurrence of positive investment and employment losses. We interpret this to mean that the anticipation of capital adjustments encourages the plant to hoard labor. In other words, the temporary decline in the marginal product of labor due to $\tau > 0$ is dwarfed by the anticipated increase once the new machinery is installed.

In a richer model, one might introduce temporary layoffs, so a plant could furlough workers while machinery was installed but recall them at no cost. It is unlikely, however, that this is important in accounting for what we observe. 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.[68] Moreover, as discussed in Section 3, there is no evidence that employment declines at investing plants are reversed at a faster rate than otherwise.

8.6 Replacement investment

In the canonical model, we assumed geometric depreciation of the plant's capital. However, it is possible that some machines are crucial to the production process. In that case, the plant will replace them even if labor demand declines. Note that, since our data do not capture pure scrappage, these required replacement investments would show up as positive *net* investment.

To investigate this quantitatively, we assume now that a portion of depreciation must be replaced *immediately*. With probability π per quarter, a fraction $\phi > 0$ of the (start-of-period) machinery stock fails and must be replaced. The remainder of the canonical model of Section 2 remains intact. To parameterize π and ϕ, we use evidence on machine scrappage and replacement, as summarized

[66] This is because disruption cost is a kind of fixed cost: *any* amount of investment incurs the discrete penalty, $(1 - \tau) \cdot F(k, n, \zeta)$. Thus, plants adjust infrequently and by "lumpy" amounts. However, this fixed-cost aspect of disruption does not, by itself, help generate investment along with employment reductions (see footnote 12).

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

[68] Note that Cooper and Haltiwanger (1993) originally applied their model to study "downtime" in vehicle assembly plants due to model-year changeovers. These are indeed typically short-lived events, in which workers are recalled within the year.

in Web Appendix F.3. This points to $\phi = 0.055$ and $\pi = 0.11$ per quarter.

The results, reported in column (4) of Table 5, appear somewhat encouraging. The model can account for about half of the periods of positive investment and negative employment growth. Moreover, the typical decline in employment in these episodes is 11 percent.

However, we remain somewhat unconvinced of this approach. What enables the model's performance is the assumption that all machine failures have to be replaced immediately. If other machines or labor could substitute for failed equipment (for even a short while), the frequency of these episodes would be much lower: plants would defer replacement until profitability recovers, which is a time when it would also hire. Eisner (1978) does indeed find that firms plan more replacements when profitability is high (see also Feldstein and Foot, 1971). This means that these results are very likely an upper bound on the implications of replacement investment for explaining why employment reductions accompany positive investment.

9 Conclusion

This paper has explored the joint dynamics of capital and employment at the plant level. We have highlighted in particular the co-occurrence of positive investment and negative employment growth at the plant level. These episodes constitute more than one-third of aggregate capital accumulation in our sample. The associated declines in employment are large and persistent.

We have explored a variety of mechanisms that might lie behind these joint dynamics. We have emphasized in particular a framework that envisions production as an aggregation of tasks. In this framework, capital can directly substitute for labor in certain tasks. This opens up a margin whereby plants endogenously adjust their capital intensity. We estimate this model against our plant-level moments—the first time, to our knowledge, any framework embodying such a mechanism has been tested on plant-level dynamics. We find that it offers a fairly comprehensive account of the joint dynamics of capital and labor adjustment.

These results suggest several avenues for continued empirical and theoretical work. For instance, more direct measurements of the uses of equipment could be instructive.[69] The Bureau of Labor Statistics (1982a,b) has provided several field studies that discuss how machinery substitutes for labor

[69]We are thankful to David Romer, who suggested this idea for a survey.

in particular tasks. However, to our knowledge, the literature lacks a comprehensive quantitative evaluation along these lines. This would offer direct evidence for (or against) the mechanism in the task framework.

In addition, it would be informative to look at the joint factor dynamics outside of manufacturing. For instance, there may be service sectors where the nature of the production process is such that workers and machines are highly complementary in the short and long run. If our interpretation of the joint factor adjustments is correct, then one should seldom see capital and labor adjust in opposite directions in these sectors. Research along these lines can thus help advance our understanding of production functions and plant-level dynamics.

Finally, while we have restricted our attention in this paper to the responses of plants to idiosyncratic shocks, a natural next step is to investigate the implications of our tasks-based production model for aggregate fluctuations. Our empirical analysis has already uncovered suggestive evidence that employment declines accompanying positive investment may be of importance for understanding business cycles: during the Asian financial crisis years of 1997-98, investing plants in our data were even *more* likely (49.7 percent of the time) to reduce employment than in 'normal times,' and the fall in employment during these episodes was larger (19.9 log points). A better understanding of joint factor dynamics is also potentially very important for understanding the effects of policy. For example, we conjecture that our tasks-based model will have very different implications for the employment response to an investment tax credit than would the workhorse Cobb-Douglas model that we have argued is strikingly inconsistent with our data. However, verifying this requires us to generalize our framework to allow for aggregate fluctuations. We plan to investigate this in future work.

A Omitted Proofs

Proof of Proposition 1. First, we verify that the plant will allocate capital to all tasks below a threshold $x \in (0, 1)$ and labor to all tasks above. Consider a plant with total capital k and workforce n. Let q_k (q_n) denote the shadow value of capital (labor). (6) then implies that the cost to the plant of producing one unit of intermediate i using capital is q_k/a_i. Likewise, the cost of producing one unit using labor is q_n/b_i. Because $b_i = 0$ for $i \leq 0$ and $a_i = 0$ for $i \geq 1$, it is clear that the plant allocates capital to tasks $i \le 0$ and labor to plants with $i \ge 1$. The plant allocates capital to task $i \in (0,1)$ if and only if

$$\frac{\xi_n q_k}{\xi_k q_n} < \frac{a_i}{b_i}.$$

Since a_i/b_i is strictly decreasing over i and satisfies $\lim_{i\to 0^+} a_i/b_i = +\infty$ and $\lim_{i\to 1^-} a_i/b_i = 0$, there is a single crossing $x \in (0,1)$ such that capital is used in tasks $i < x$ and labor in tasks $i > x$.

Now for a given threshold x, the plant allocates its stocks k and n across the continuum of tasks in order to maximize total output. By (6), this problem can also be cast in terms of choosing the outputs of intermediates $\{y_i\}_{i=0}^{1}$. We thus envision the plant selecting $\{y_i\}_{i=0}^{1}$ to maximize[70]

$$Y = \left[\int_{-d}^{1+d} y_i^{\frac{\varphi-1}{\varphi}}\,\mathrm{d}i\right]^{\frac{\varphi}{\varphi-1}} \tag{A.1}$$

subject to the resource constraints,

$$k = \int_{-d}^{x} k_i\,\mathrm{d}i, \qquad n = \int_{x}^{1+d} n_i\,\mathrm{d}i, \tag{A.2}$$

and (6). The first-order conditions are:

$$i < x: \qquad (y_i/Y)^{-1/\varphi} = q_k/\left(\xi_k a_i\right), \tag{A.3}$$

$$i \ge x: \qquad (y_i/Y)^{-1/\varphi} = q_n/\left(\xi_n b_i\right). \tag{A.4}$$

To simplify, recall that on tasks $i \ge x$, we have $y_i = \xi_n b_i n_i$. Using $n_i = y_i/\left[\xi_n b_i\right]$ in (A.2) and substituting for y_i using equation (A.4) yields

$$n = Y q_n^{-\varphi} \xi_n^{\varphi-1} \int_{x}^{1+d} b_i^{\varphi-1}\,\mathrm{d}i.$$

Now using this to replace $Y q_n^{-\varphi}$ in (A.4) gives

$$i \ge x: \qquad y_i = \xi_n n \frac{b_i^{\varphi}}{B(x)}, \tag{A.5}$$

where $B(x) \equiv \int_x^{1+d} b_i^{\varphi-1}\,\mathrm{d}i$. Next, on tasks $i < x$, $y_i = \xi_k a_i k_i$. Replacing k_i in (A.2) with $y_i/\left[\xi_k a_i\right]$;

[70] Notice that conditional on k and n, maximizing revenue is equivalent to maximizing output.

substituting for y_i from (A.3); and using (A.3) again to remove $Yq_k^{-\varphi}$ yields

$$i < x: \qquad y_i = \xi_k k \frac{a_i^{\varphi}}{A(x)}, \tag{A.6}$$

where $A(x) \equiv \int_{-d}^{x} a_i^{\varphi-1}\,\mathrm{d}i$. Now substituting (A.5) and (A.6) into (A.1) completes the proof by showing that the production function takes the form

$$F(k,n,x) = \left[A(x)^{1/\varphi}(\xi_k k)^{\frac{\varphi-1}{\varphi}} + B(x)^{1/\varphi}(\xi_n n)^{\frac{\varphi-1}{\varphi}}\right]^{\frac{\varphi}{\varphi-1}}.$$

Proof of Proposition 2. Differentiating (7) with respect to k and n and using the envelope theorem, we have

$$F_k^* = \left[A(x)^{1/\varphi} + B(x)^{1/\varphi}\left(\frac{\xi_n n}{\xi_k k}\right)^{\frac{\varphi-1}{\varphi}}\right]^{\frac{1}{\varphi-1}} \times A(x)^{1/\varphi}\,\xi_k$$

and

$$F_n^* = \left[A(x)^{1/\varphi} + B(x)^{1/\varphi}\left(\frac{\xi_n n}{\xi_k k}\right)^{\frac{\varphi-1}{\varphi}}\right]^{\frac{1}{\varphi-1}} \times \left(\frac{\xi_k k}{\xi_n n}\right)^{1/\varphi} B(x)^{1/\varphi}\,\xi_n.$$

It follows that

$$\ln\frac{F_k^*}{F_n^*} = \frac{1}{\varphi}\ln\frac{A(x)}{B(x)} - \frac{1}{\varphi}\ln\frac{k}{n} + \frac{\varphi-1}{\varphi}\ln\frac{\xi_k}{\xi_n},$$

and

$$-\frac{\partial\ln\left[F_k^*/F_n^*\right]}{\partial\ln\left[k/n\right]} = \frac{1}{\varphi}\left\{1 - \frac{\partial\ln\left[A(x)/B(x)\right]}{\partial\ln\left[k/n\right]}\right\}.$$

Therefore, $-\left[\partial\ln\left(F_k^*/F_n^*\right)/\partial\ln\left(k/n\right)\right]^{-1} > \varphi$ if $\partial\ln\left[A(x)/B(x)\right]/\partial\ln\left[k/n\right] > 0$. This can be verified using the first-order condition for x. The latter implies that

$$\ln\frac{A(x)}{B(x)} + \frac{\varphi}{1-\varphi}\ln\left[-\frac{A'(x)}{B'(x)}\right] = \ln\frac{\xi_k k}{\xi_n n},$$

or, after using $A'(x) = a(x)^{\varphi-1}$ and $B'(x) = -b(x)^{\varphi-1}$,

$$\ln\frac{A(x)}{B(x)} - \varphi\ln\left[\frac{a(x)}{b(x)}\right] = \ln\frac{\xi_k k}{\xi_n n}.$$

If $\ln(k/n)$ increases, the left side must rise. Since $A(x)/B(x)$ is increasing in x and $a(x)/b(x)$ is decreasing, the left side can only rise if x rises. Thus, $\partial\ln\left[A(x)/B(x)\right]/\partial\ln\left[k/n\right] > 0$.

Proof of Proposition 3. First suppose $\psi > \max\{\varphi, 1\}$. Notice that because $\psi > 1$, we have that a_i is strictly decreasing in i and b_i is strictly increasing in i. To verify (9), calculate

$$A(x) \equiv \int_0^x a_i^{\varphi-1}\,\mathrm{d}i = \chi^{\varphi-1}\int_0^x i^{-\frac{\varphi-1}{\psi-1}}\,\mathrm{d}i = \chi^{\varphi-1}\frac{\psi-1}{\psi-\varphi}x^{\frac{\psi-\varphi}{\psi-1}} = x^{\varphi\Gamma};$$
$$B(x) \equiv \int_x^1 b_i^{\varphi-1}\,\mathrm{d}i = \chi^{\varphi-1}\int_x^1 (1-i)^{-\frac{\varphi-1}{\psi-1}}\,\mathrm{d}i = \chi^{\varphi-1}\frac{\psi-1}{\psi-\varphi}(1-x)^{\frac{\psi-\varphi}{\psi-1}} = (1-x)^{\varphi\Gamma}.$$

In making this calculation, we need $\varphi\Gamma \equiv (\psi-\varphi)/(\psi-1)$ to be strictly positive so that the integrals defining $A(x)$ and $B(x)$ converge; this is guaranteed by the assumption that $\psi > \max\{\varphi, 1\}$.

Now choose x to maximize (9). The derivative of the production function with respect to x is

$$\frac{\partial}{\partial x}F(k,n,x) = \Gamma\frac{\varphi}{\varphi-1}F(k,n,x)^{\frac{1}{\varphi}}D(k,n,x) \tag{A.7}$$

where $D(k,n,x) = x^{\Gamma-1}(\xi_k k)^{(\varphi-1)/\varphi} - (1-x)^{\Gamma-1}(\xi_n n)^{(\varphi-1)/\varphi}$. The first-order condition for x is therefore satisfied iff $D(k,n,x) = 0$, or equivalently, if $x/(1-x) = [(\xi_k k)/(\xi_n n)]^{(\psi-1)/\psi}$, that is

$$x = x^* \equiv \frac{(\xi_k k)^{\frac{\psi-1}{\psi}}}{(\xi_k k)^{\frac{\psi-1}{\psi}} + (\xi_n n)^{\frac{\psi-1}{\psi}}}, \tag{A.8}$$

which is necessarily positive and strictly less than 1. To verify that this solution is a maximum, evaluate the second derivative at the optimal x:

$$\left.\frac{\partial^2}{\partial x^2}F(k,n,x)\right|_{x=x^*} = \Gamma(\Gamma-1)\frac{\varphi}{\varphi-1}F(k,n,x^*)^{\frac{1}{\varphi}}\left[x^{*\Gamma-2}(\xi_k k)^{\frac{\varphi-1}{\varphi}} + (1-x^*)^{\Gamma-2}(\xi_n n)^{\frac{\varphi-1}{\varphi}}\right].$$

This expression has the same sign as $\Gamma(\Gamma-1)\varphi/(\varphi-1) = -(\psi-\varphi)\psi/[\varphi(\psi-1)^2]$. The assumption that $\psi > \varphi$ ensures that this is negative. Finally, substitute from (A.8) into (9) to establish (11).

Now suppose $\varphi < \psi < 1$. Since $\psi < 1$, we again have that a_i is strictly decreasing and b_i strictly increasing on $(0,1)$. Substitute for $\{a_i, b_i\}$ to calculate

$$A(x) = \int_{-1}^x a_i^{\varphi-1}\,\mathrm{d}i = \int_{-1}^0 \mathrm{d}i + \chi^{\varphi-1}\int_0^x (1-i)^{\frac{\varphi-1}{1-\psi}}\,\mathrm{d}i = 1 + \chi^{\varphi-1}\frac{1-\psi}{\varphi-\psi}\left[(1-x)^{\frac{\varphi-\psi}{1-\psi}} - 1\right] = (1-x)^{\varphi\Gamma};$$
$$B(x) = \int_x^2 b_i^{\varphi-1}\,\mathrm{d}i = \chi^{\varphi-1}\int_x^1 i^{\frac{\varphi-1}{1-\psi}}\,\mathrm{d}i + \int_1^2 \mathrm{d}i = \chi^{\varphi-1}\frac{1-\psi}{\varphi-\psi}\left[1 - x^{\frac{\varphi-\psi}{1-\psi}}\right] + 1 = x^{\varphi\Gamma}.$$

This establishes (10). The derivative of $F(k,n,x)$ with respect to x is again given by (A.7) provided we redefine $D(k,n,x) = x^{\Gamma-1}(\xi_n n)^{(\varphi-1)/\varphi} - (1-x)^{\Gamma-1}(\xi_k k)^{(\varphi-1)/\varphi}$. The first-order condition for x is thus satisfied if $x/(1-x) = [(\xi_n n)\,/\,(\xi_k k)]^{(\psi-1)/\psi}$, that is $x = 1 - x^*$ where x^* is as in (A.8). It is again straightforward to verify that this is a maximum using the second-order condition. Finally, substitute for x into (10) to establish (11).

Proof of Proposition 4. We consider here only the case where $\psi > \max\{\varphi, 1\}$; a very similar argument applies in the case where $\varphi < \psi < 1$, and is given in Web Appendix C.

The firm's problem is (12) where the production function takes the form (9). First consider the barriers for capital. If $i \neq 0$, the form of the capital adjustment cost (4) implies the first-order condition for the choice of k takes the form

$$\zeta^{\frac{1}{\epsilon}} \left[x^{\Gamma}(\xi_k k)^{\frac{\varphi-1}{\varphi}} + (1-x)^{\Gamma}(\xi_n n)^{\frac{\varphi-1}{\varphi}}\right]^{\frac{\epsilon-\varphi}{\epsilon(\varphi-1)}} x^{\Gamma} \xi_k^{\frac{\varphi-1}{\varphi}} k^{-\frac{1}{\varphi}} = r c_k \tag{A.9}$$

where $r \equiv 1 - \varrho$ and where $c_k = 1$ if $i > 0$ and $c_k = p_s$ if $i < 0$. Write $\underline{k}(n,x)$ for the investing barrier given n and x, that is, the value of k that satisfies (A.9) for $c_k = 1$. Write $\eta_{\underline{k},x} \equiv \partial \ln \underline{k}(n,x)/\partial \ln x$, the elasticity of $\underline{k}(n,x)$ with respect to x. Log-differentiating (A.9) with respect to x implies that

$$\frac{\epsilon-\varphi}{\epsilon(\varphi-1)}\left[\Gamma s - \Gamma\frac{x}{1-x}(1-s) + \frac{\varphi-1}{\varphi} s\eta_{\underline{k},x}\right] + \Gamma - \frac{1}{\varphi}\eta_{\underline{k},x} = 0 \tag{A.10}$$

where

$$s = \frac{x^{\Gamma}(\xi_k k)^{\frac{\varphi-1}{\varphi}}}{x^{\Gamma}(\xi_k k)^{\frac{\varphi-1}{\varphi}} + (1-x)^{\Gamma}(\xi_n n)^{\frac{\varphi-1}{\varphi}}}. \tag{A.11}$$

(It is straightforward to verify that s is the capital cost share $rk/(rk+wn)$ in the frictionless case with $p_s = 1$ and $c^+ = c^- = 0$, that is, $c_k \equiv 1$ and $c_n \equiv 0$, but we do not need to use this fact here.) Conditional on k and n, the first-order condition for the optimal choice of x implies that $s/x = (1-s)/(1-x)$. Use this to simplify the parenthesis on the left side of (A.10), multiply the last term on the left side of the equation by $s + (1-s) = 1$, and rearrange to obtain

$$\left[\left(\frac{1}{\varphi} - \frac{\epsilon-\varphi}{\epsilon\varphi}\right)s + \frac{1}{\varphi}(1-s)\right]\eta_{\underline{k},x} = \Gamma,$$

that is,

$$\eta_{\underline{k},x} = \frac{\Gamma}{\frac{1}{\epsilon}s + \frac{1}{\varphi}(1-s)} > 0.$$

An essentially identical calculation shows that $\eta_{\bar{k},x}$, the elasticity of the disinvesting barrier with respect to x, takes the same value.

Now consider the barriers for labor. If $\Delta n \neq 0$, the first-order condition for labor takes the form

$$\zeta^{\frac{1}{\epsilon}} \left[x^{\Gamma} (\xi_k k)^{\frac{\varphi-1}{\varphi}} + (1-x)^{\Gamma} (\xi_n n)^{\frac{\varphi-1}{\varphi}} \right]^{\frac{\epsilon-\varphi}{\epsilon(\varphi-1)}} (1-x)^{\Gamma} \xi_n^{\frac{\varphi-1}{\varphi}} n^{-\frac{1}{\varphi}} = w + rc_n, \tag{A.12}$$

where $c_n = c^+$ if $\Delta n > 0$ and $c_n = -c^-$ if $\Delta n < 0$. Again, log-differentiate (A.12) with respect to x, holding k constant, to obtain that the elasticity of the hiring barrier, $\eta_{\underline{n},x}$, satisfies

$$\frac{\epsilon-\varphi}{\epsilon(\varphi-1)} \left[\Gamma s - \Gamma \frac{x}{1-x}(1-s) + \frac{\varphi-1}{\varphi}(1-s)\eta_{\underline{n},x} \right] - \Gamma \frac{x}{1-x} - \frac{1}{\varphi}\eta_{\underline{n},x} = 0.$$

Substitute $s/x = (1-s)/(1-x)$ and simplify as before to see that

$$\eta_{\underline{n},x} = -\frac{\Gamma}{\frac{1}{\varphi}s + \frac{1}{\epsilon}(1-s)} \frac{x}{1-x} < 0.$$

Again, $\eta_{\bar{n},x}$ takes the same value.

References

Abel, Andrew and Olivier J. Blanchard. 1988. "Investment and Sales: Some Empirical Evidence." In *Dynamic Economic Modelling: Proceedings of the Third International Symposium in Economic Theory & Econometrics*, edited by William A. Barnett, Ernst Berndt, and Halbert White. New York: Cambridge University Press, 269–296.

Abel, Andrew and Janice Eberly. 1996. "Optimal Investment with Costly Reversibility." *Review of Economic Studies* 63 (4):581–594.

Acemoglu, Daron. 2010. "When Does Labor Scarcity Encourage Innovation?" *Journal of Political Economy* 118 (6):1037–1078.

Acemoglu, Daron and David Autor. 2011. "Skills, Tasks, and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, vol. 4B, edited by David Card and Orley Ashenfelter, chap. 12. Amsterdam: Elsevier, 1043–1171.

Acemoglu, Daron and William B. Hawkins. 2014. "Search with Multi-Worker Firms." *Theoretical Economics* 9 (3):583–628.

Acemoglu, Daron and Pascual Restrepo. 2014. "The Race Between Machine and Man: Implications of Technology for Growth, Factor Shares and Unemployment." Unpublished.

Acemoglu, Daron and Fabrizio Zilibotti. 2001. "Productivity Differences." *Quarterly Journal of Economics* 116 (2):563–606.

Altonji, Joseph G and Lewis M. Segal. 1996. "Small-Sample Bias in GMM Estimation of Covariance Structures." *Journal of Business and Economic Statistics* 14 (3):353–366.

Antony, Jürgen. 2009. "A dual elasticity of substitution production function with an application to cross-country inequality." *Economics Letters* 102 (1):10–12.

———. 2014. "Technical change and the elasticity of factor substitution." Beiträge der Hochschule Pforzheim 147, Pforzheim, Germany.

Arellano, Cristina, Yan Bai, and Patrick Kehoe. 2012. "Financial Frictions and Fluctuations in Volatility." Staff Report 466, Federal Reserve Bank of Minneapolis, Minneapolis, MN.

Basu, Susanto, John G. Fernald, and Miles S. Kimball. 2006. "Are Technology Improvements Contractionary?" *American Economic Review* 96 (5):1418–1448.

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

Berman, Eli, John Bound, and Steve Machin. 1998. "Implications of Skill-Biased Technological Change: International Evidence." *Quarterly Journal of Economics* 113 (4):1245–1279.

Berman, Eli and Steve Machin. 2000. "Skill-Biased Technology Transfer Around the World." *Oxford Review of Economic Policy* 16 (3):12–22.

Bloom, Nicholas. 2009. "The impact of uncertainty shocks." *Econometrica* 77 (3):623–685.

Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. 2012. "Really Uncertain Business Cycles." Working Paper 18245, NBER, Cambridge, MA.

Bureau of Labor Statistics. 1982a. "The Impact of Technology on Labor in Four Industries: Printing and publishing, Water transportation, Copper ore mining, Fabricated structural metal, Intercity trucking." Bulletin 2137, Bureau of Labor Statistics, Washington, DC.

———. 1982b. "Technology and Labor in Four Industries: Meat Products, Foundries, Metalworking Machinery, Electrical and Electronic Equipment." Bulletin 2104, Bureau of Labor Statistics, Washington, DC.

Caballero, Ricardo J. and Eduardo M. R. A. Engel. 1999. "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach." *Econometrica* 67 (4):783–826.

Caballero, Ricardo J., Eduardo M. R. A. Engel, and John C. Haltiwanger. 1995. "Plant-Level Adjustment and Aggregate Investment Dynamics." *Brookings Papers on Economic Activity* 26 (2):1–54.

———. 1997. "Aggregate Employment Dynamics: Building from Microeconomic Evidence." *American Economic Review* 87 (1):115–137.

Caballero, Ricardo J. and Mohamad L. Hammour. 1998. “Jobless growth: appropriability, factor substitution, and unemployment.” *Carnegie-Rochester Conference Series on Public Policy* 48:51–94.

Calvo, Guillermo A. 1983. “Staggered prices in a utility-maximizing framework.” *Journal of Monetary Economics* 12 (3):383–398.

Caselli, Francesco and Wilbur John Coleman, II. 2006. “The World Technology Frontier.” *American Economic Review* 96 (3):499–522.

Chirinko, Robert S. 2008. “σ: The long and short of it.” *Journal of Macroeconomics* 30 (2):671–686.

Chirinko, Robert S., Steven M. Fazzari, and Andrew P. Meyer. 2011. “A New Approach to Estimating Production Function Parameters: The Elusive Capital-Labor Substitution Elasticity.” *Journal of Business and Economic Statistics* 29 (4):587–594.

Cooper, Russell W. and John C. Haltiwanger. 1993. “The Aggregate Implications of Machine Replacement: Theory and Evidence.” *American Economic Review* 83 (3):360–382.

———. 2000. “On the Nature of Capital Adjustment Costs.” Working Paper 7925, NBER, Cambridge, MA.

———. 2006. “On the Nature of Capital Adjustment Costs.” *Review of Economic Studies* 73 (3):611–633.

Cooper, Russell W., John C. Haltiwanger, and Laura Power. 1999. “Machine Replacement and the Business Cycle: Lumps and Bumps.” *American Economic Review* 89 (4):921–946.

Cooper, Russell W., John C. Haltiwanger, and Jonathan L. Willis. 2004. “Dynamics of Labor Demand: Evidence from Plant-level Observations and Aggregate Implications.” Working Paper 10297, NBER, Cambridge, MA.

———. 2007. “Search frictions: Matching aggregate and establishment observations.” *Journal of Monetary Economics* 54 (Supplement 1):56–78.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2006. “The Flow Approach to Labor Markets: New Data Sources and Micro-Macro Links.” *Journal of Economic Perspectives* 20 (3):3–26.

———. 2013. “The Establishment-Level Behavior of Vacancies and Hiring.” *Quarterly Journal of Economics* 128 (2):581–622.

Davis, Steven J. and John C. Haltiwanger. 1990. “Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications.” *NBER Macroeconomics Annual* 5:123–168.

———. 1992. “Gross Job Creation, Gross Job Destruction, and Employment Reallocation.” *Quarterly Journal of Economics* 107 (3):819–863.

Diamond, Peter, Daniel McFadden, and Miguel Rodriguez. 1978. “Measurement of the Elasticity of Substitution and Bias of Technical Change.” In *Production Economics: A Dual Approach to Theory and Applications Volume II: Applications of the Theory of Production*, *Contributions to Economic Analysis*, vol. 111, edited by Melvyn Fuss and Daniel McFadden, chap. IV.2. Amsterdam: North-Holland, 125–147.

Dixit, Avinash. 1997. "Investment and Employment Dynamics in the Short Run and the Long Run." *Oxford Economic Papers* 49 (1):1–20.

Doms, Mark and Timothy Dunne. 1994. "Capital Adjustment Patterns in Manufacturing Plants." Working Paper 94-11, Center for Economic Studies, U.S. Census Bureau, Washington, DC.

———. 1998. "Capital adjustment patterns in manufacturing plants." *Review of Economic Dynamics* 1 (2):409–429.

Eberly, Janice and Jan van Mieghem. 1997. "Multi-factor Dynamic Investment Under Uncertainty." *Journal of Economic Theory* 75 (2):345–387.

Eisner, Robert. 1978. "Components of Capital Expenditures – Replacement and Modernization versus Expansion." In *Factors in Business Investment*, chap. 8. Cambridge, MA: NBER, 175–188.

Elsby, Michael and Ryan Michaels. 2013. "Marginal Jobs, Heterogeneous Firms, and Unemployment Flows." *American Economic Journal: Macroeconomics* 5 (1):1–48.

Eslava, Marcela, John C. Haltiwanger, Adriana Kugler, and Maurice Kugler. 2010. "Factor Adjustment After Deregulation: Panel Evidence from Colombian Plants." *Review of Economics and Statistics* 92 (2):378–391.

Feldstein, Martin S. and David K. Foot. 1971. "The Other Half of Gross Investment: Replacement and Modernization Expenditures." *Review of Economics and Statistics* 53 (1):49–58.

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2008. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review* 98 (1):394–425.

Goldin, Claudia and Lawrence F. Katz. 1998. "The Origins Of Technology-Skill Complementarity." *Quarterly Journal of Economics* 113 (3):693–732.

Gourio, François and Anil K. Kashyap. 2007. "Investment spikes: New facts and a general equilibrium exploration." *Journal of Monetary Economics* 54 (Supplement):1–22.

Griliches, Zvi. 1969. "Capital-Skill Complementarity." *Review of Economics and Statistics* 51 (4):465–468.

Habakkuk, H. J. 1962. *American and British Technology in the Nineteenth Century*. London: Cambridge University Press.

Hall, Robert E. 2004. "Measuring Factor Adjustment Costs." *Quarterly Journal of Economics* 119 (3):899–927.

Houthakker, H. S. 1955-56. "The Pareto Distribution and the Cobb-Douglas Production Function in Activity Analysis." *Review of Economic Studies* 23 (1):27–31.

Johansen, Leif. 1959. "Substitution versus Fixed Production Coefficients in the Theory of Economic Growth: A Synthesis." *Econometrica* 27 (2):157–176.

Jones, Charles I. 2005. "The Shape of the Production Function and the Direction of Technical Change." *Quarterly Journal of Economics* 120 (2):517–549.

Jovanovic, Boyan and Dmitriy Stolyarov. 2000. "Optimal Adoption of Complementary Technologies." *American Economic Review* 90 (1):15–29.

Kaas, Leo and Philipp Kircher. 2011. “Efficient Firm Dynamics in a Frictional Labor Market.” Unpublished.

Karabarbounis, Loukas and Brent Neiman. 2014. “The Global Decline of the Labor Share.” *Quarterly Journal of Economics* 129 (1):61–103.

Khan, Aubhik and Julia K. Thomas. 2008. “Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics.” *Econometrica* 76 (2):395–436.

———. 2013. “Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity.” *Journal of Political Economy* 121 (6):1055–1107.

Klump, Rainer, Peter McAdam, and Alpo Willman. 2007. “Factor Substitution and Factor-Augmenting Technical Progress in the United States: A Normalized Supply-Side System Approach.” *Review of Economics and Statistics* 89 (1):183–192.

Krusell, Per, Lee E. Ohanian, José-Víctor Ríos-Rull, and Giovanni L. Violante. 2000. “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis.” *Econometrica* 68 (5):1029–1053.

Léon-Ledesma, Miguel and Mathan Satchi. 2011. “The choice of capital intensity and balanced growth.” Unpublished.

Léon-Ledesma, Miguel A., Peter McAdam, and Alpo Willman. 2010. “Identifying the Elasticity of Substitution with Biased Technical Change.” *American Economic Review* 100 (4):1330–1357.

Letterie, Wilko, Gerard Pfann, and J. Michael Polder. 2004. “Factor adjustment spikes and interrelation: an empirical investigation.” *Economics Letters* 85 (2):145–150.

Merz, Monika and Eran Yashiv. 2007. “Labor and the Market Value of a Firm.” *American Economic Review* 97 (4):1419–1431.

Moscarini, Giuseppe and Fabien Postel-Vinay. 2009. “The Timing of Labor Market Expansions: New Facts and a New Hypothesis.” In *NBER Macroeconomics Annual*, vol. 23, edited by K. Daron Acemoglu, Kenneth Rogoff, and Michael Woodford. NBER, 1–51.

———. 2012. “The Contribution of Large and Small Employers to Job Creation in Times of High and Low Unemployment.” *American Economic Review* 102 (6):2509–2539.

Nadiri, M. Ishag and Sherwin Rosen. 1969. “Interrelated Factor Demand Functions.” *American Economic Review* 59 (4):457–471.

Nakamura, Emi and Jon Steinsson. 2008. “Five Facts about Prices: A Reevaluation of Menu Cost Models.” *Quarterly Journal of Economics* 123 (4):1415–1464.

Nakamura, Hideki. 2010. “Factor Substitution, Mechanization, and Economic Growth.” *Japanese Economic Review* 61 (2):266–281.

Nakamura, Hideki and Masakatsu Nakamura. 2008. “Constant-Elasticity-of-Substitution Production Function.” *Macroeconomic Dynamics* 12 (5):694–701.

Nilsen, Oivind and Fabio Schiantarelli. 2003. “Zeroes and Lumps in Investment: Empirical Evidence on Irreversibilities and Nonconvexities.” *Review of Economics and Statistics* 85 (4):1021–1043.

Oberfield, Ezra and Devesh Raval. 2014. “Micro Data and Macro Technology.” Unpublished.

Oostendorp, Remco. 2012. “The Occupational Wages around the World (OWW) Database: Update for 1983-2008.” Unpublished.

Polder, Michael and Sher Verick. 2004. “Dynamics of Labour and Capital Adjustment - A Comparison of Germany and the Netherlands.” Discussion Paper 1212, IZA, Bonn, Germany.

Raval, Devesh. 2014. “The Micro Elasticity of Substitution and Non Neutral Technology.” Unpublished.

Reiter, Michael, Tommy Sveen, and Lutz Weinke. 2012. “Lumpy Investment and State-Dependent Pricing in General Equilibrium.” Unpublished.

Rossana, Robert J. 1990. “Interrelated Demands for Buffer Stocks and Productive Inputs: Estimates for Two-Digit Manufacturing Industries.” *Review of Economics and Statistics* 72 (1):19–29.

Roys, Nicolas. 2014. “Persistence of Shocks and the Reallocation of Labor.” Unpublished.

Sakellaris, Plutarchos. 2001. “Patterns of Plant Adjustment.” Finance and Economics Discussion Series 5, Federal Reserve Board, Washington, DC.

———. 2004. “Patterns of Plant Adjustment.” *Journal of Monetary Economics* 51 (2):425–450.

Samuelson, Paul Anthony. 1947. *Foundations of Economic Analysis*. Cambridge, MA: Harvard University Press.

Schaal, Edouard. 2012. “Uncertainty, Productivity and Unemployment during the Great Recession.” Unpublished.

Shapiro, Matthew D. 1986. “The Dynamic Demand for Capital and Labor.” *Quarterly Journal of Economics* 101 (3):513–542.

Smith, Eric. 1999. “Search, Concave Production and Optimal Firm Size.” *Review of Economic Dynamics* 2 (2):456–471.

Stole, Lars A. and Jeffrey Zwiebel. 1996. “Intra-firm Bargaining under Non-binding Contracts.” *Review of Economic Studies* 63 (3):375–410.

Thomas, Julia K. 2002. “Is Lumpy Investment Relevant for the Business Cycle?” *Journal of Political Economy* 110:508–534.

Veracierto, Marcelo. 2002. “Plant Level Irreversible Investment and Equilibrium Business Cycles.” *American Economic Review* 92 (1):181–197.

———. 2009. “Establishments Dynamics, Vacancies and Unemployment: A Neoclassical Synthesis.” Working Paper 2009-14, Federal Reserve Bank of Chicago, Chicago, IL.

Zeira, Joseph. 1998. “Workers, Machines and Economic Growth.” *Quarterly Journal of Economics* 113 (4):1091–1113.

———. 2006. “Machines as Engines of Growth.” Unpublished.