NBER WORKING PAPER SERIES

CAPITAL REALLOCATION AND AGGREGATE PRODUCTIVITY

Russell W. Cooper
Immo Schott

Working Paper 19715
http://www.nber.org/papers/w19715

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2013

Thanks to Dean Corbae for lengthy discussions on a related project. The first author thanks the NSF under grant #0819682 for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Capital Reallocation and Aggregate Productivity
Russell W. Cooper and Immo Schott
NBER Working Paper No. 19715
December 2013
JEL No. D24,E2,E22,E32

ABSTRACT

This paper studies the productivity implications of cyclical reallocation. Frictions in the reallocation process are a source of factor misallocation. Cyclical movements in these frictions lead to variations in the degree of reallocation and thus in productivity. These frictions also impact the capital accumulation decision. The effects are quantitatively important in the presence of fluctuations in adjustment frictions and/or the cross sectional variation of profitability shocks.

Russell W. Cooper
Department of Economics
The Pennsylvania State University
611 Kern
State College, PA 16802
and NBER
russellcoop@gmail.com

Immo Schott
European University Institute
Via della Piazzuola 43
50133 Firenze, Italy
immo.schott@eui.eu

Capital Reallocation and Aggregate Productivity*

Russell W. Cooper† and Immo Schott‡

December 4, 2013

Abstract

This paper studies the productivity implications of cyclical reallocation. Frictions in the reallocation process are a source of factor misallocation. Cyclical movements in these frictions lead to variations in the degree of reallocation and thus in productivity. These frictions also impact the capital accumulation decision. The effects are quantitatively important in the presence of fluctuations in adjustment frictions and/or the cross sectional variation of profitability shocks.

1 Motivation

Frictions in the reallocation of capital and labor are important for understanding aggregate productivity. With heterogenous plants and heterogenous inputs of capital and labor, the assignment of inputs to firms impacts directly on aggregate productivity. Frictions in the reallocation process thus lead to the misallocation of factors of production (relative to a frictionless benchmark). This point lies at the heart of the analysis of productivity across countries in Hsieh and Klenow (2007), Bartelsman, Haltiwanger, and Scarpetta (2006) and Restuccia and Rogerson (2007).[1]

We consider the cyclical dimension of reallocation in the presence of capital adjustment costs. In an important empirical contribution, Eisfeldt and Rampini (2006) show that capital reallocation is procyclical and that the cross-sectional productivity dispersion behaves countercyclically.[2] This not only underlines the importance of heterogeneity in the produc-

*Thanks to Dean Corbae for lengthy discussions on a related project. The first author thanks the NSF under grant #0819682 for financial support.

†Department of Economics, the Pennsylvania State University and NBER, russellcoop@gmail.com

‡Department of Economics, European University Institute, Immo.Schott@eui.eu

[1] More specific differences with these and other studies are discussed below.

[2] Eisfeldt and Rampini (2006) use several measures of productivity dispersion: dispersion in firm level Tobins q, dispersion in firm level investment rates, dispersion in total factor productivity growth rates, and dispersion in capacity utilization. See also Kehrig (2011) for an empirical investigation of countercyclical productivity dispersion.

tion sector but also suggests that frictions in the adjustment to capital may produce cyclical effects on output over the business cycle.

Not properly taking cross-sectional heterogeneity into account will also lead to a mis-measurement of total factor productivity (TFP). We are interested in the cyclical component of the output loss resulting from frictions in the adjustment process which will be reflected in a mis-measured TFP. This relates to the question how micro-frictions like physical adjustment costs translate into aggregate outcomes.

We find that if the only shocks in the economy are to aggregate TFP, then the productivity loss from costly reallocation has no cyclical element. This is consistent with results on the aggregate implications of lumpy investment, as in Thomas (2002), Khan and Thomas (2003) and Gourio and Kashyap (2007). If an aggregate model behaves as if there were no non-convexities at the plant-level, then the distortions in the allocation of capital across plants with different productivities will matter only for aggregate *levels*. As a result, the distribution over plants' capital stock and idiosyncratic productivity can be extremely well approximated by its first moment.

However, shocks to either adjustment costs or to the cross sectional distribution of idiosyncratic shocks create cyclical movements in reallocation and productivity. Cross-sectional heterogeneity plays an important role for shaping aggregate dynamics. In the presence of those shocks, reallocation is correlated with measured aggregate productivity. The cross-sectional joint distribution over plants' capital stock and idiosyncratic productivity is a slow-moving object in this environment. Because changes to the variance of idiosyncratic shocks do not effect all moments of this distribution in the same way, tracking its evolution only by its first moment is insufficient: higher order moments are needed to characterize the outcome of the planner's problem, in particular the covariance of the cross-sectional distribution between plants' capital stocks and profitability.

Importantly these features of our model are all related. The fact that the covariance matters as a moment for determining the optimal allocation is indicative of the significance of reallocation effects. If this covariance did not matter for describing optimal allocations - for example because it is constant over time or perfectly correlated with the mean - then it could not have a cyclical effect on aggregate output. Thus the covariance that matters from the perspective of the Krusell and Smith (1998) approach is precisely the moment that reflects gains to capital reallocation.

Studies following Krusell and Smith (1998) routinely find that only first moments of distributions are needed to summarize cross sectional distributions. In our economy, the covariance of the cross sectional distribution between a plant's capital and its profitability is needed in the state space of the problem. If there are exclusively TFP shocks, then this covariance is only present as a constant. When there are shocks either to the capital adjustment process or to the cross sectional distribution, this covariance evolves in response to these shocks. In the presence of such shocks the approximate solution to the planner's problem using only average capital fails: the solution requires higher order moments.

2 Frictionless Economy

To fix basic ideas and notation, start with an economy with heterogeneity and no frictions. The planner maximizes

$$V(A,K) = max_{K',k(\varepsilon)}\ u(c) + \beta E_{A'|A} V(A',K') \tag{1}$$

for all (A,K). The constraints are

$$c + K' = y + (1-\delta)K, \tag{2}$$

$$\int_{\varepsilon} k(\varepsilon) f(\varepsilon) d\varepsilon = K, \tag{3}$$

$$y = A \int_{\varepsilon} \varepsilon k(\varepsilon)^{\alpha} f(\varepsilon) d(\varepsilon). \tag{4}$$

The objective function is the lifetime utility of the representative household. The state vector has two elements: A is aggregate TFP and K is the aggregate stock of capital. There is a distribution of plant specific productivity shocks, $f(\varepsilon)$ which is fixed and hence omitted from the state vector. Further, as there is costless reallocation of capital, the joint distribution of plant-specific capital and ε is not part of the state vector.

There are two controls in (1). The first is the choice of aggregate capital for the next period. The second is the assignment function, $k(\varepsilon)$, which allocates the given stock of capital across the production sites, indexed by their current productivity.

At the beginning of the period, A as well as the idiosyncratic productivity shocks ε realize. While aggregate capital K requires one period time-to-build, the reallocation of existing capital takes place instantaneously and is given by $k(\varepsilon)$.

The resource constraint for the accumulation of aggregate capital is given in (2). The constraint for the allocation of capital across production sites in given in (3), where $f(\cdot)$ is the current distribution function for ε. From (4), total output, y, is the sum of the output across production sites. The production function at any site is $y(k,\varepsilon) = A\varepsilon k^{\alpha}$ where k is the capital used at the site with productivity ε. The idiosyncratic productivity ε is persistent, parameterized by $\rho_{\varepsilon} \in [0,1]$. We assume $\alpha < 1$ as in Lucas (1978). In this frictionless environment, a plants' optimal capital stock is entirely determined by ε.

From (4), total output is then the integration over production sites under the allocation rule $k(\varepsilon)$ given the cross sectional distribution of productivity, $f(\cdot)$. The assumption of diminishing returns to scale, $\alpha < 1$, implies that the allocation of capital across production sites is non-trivial. There are gains to allocating capital to high productivity sites but there are also gains, due to $\alpha < 1$, from spreading capital out across the production sites.

2.1 Optimal Choices

Within a period, the condition for the optimal allocation of capital across production sites is given by $\alpha A \varepsilon k(\varepsilon)^{\alpha-1} = \eta$ for all ε, where η is the multiplier on (3). This condition is intuitive: absent frictions, the optimal allocation equates the marginal product of capital across production sites.

Working with this condition,

$$k(\varepsilon) = {\frac{\eta}{\alpha A \varepsilon}}^{\frac{1}{\alpha-1}}. \tag{5}$$

Using (3),

$$\eta = A\alpha K^{\alpha-1} \left(\int_\varepsilon \varepsilon^{\frac{1}{1-\alpha}} f(\varepsilon) d\varepsilon \right)^{1-\alpha}. \tag{6}$$

The multiplier is the standard marginal product on an additional unit of capital times the effect of the ε distribution on productivity.

Putting these two conditions together,

$$k(\varepsilon) = K \frac{\varepsilon^{\frac{1}{1-\alpha}}}{\int_\varepsilon \varepsilon^{\frac{1}{1-\alpha}} f(\varepsilon) d\varepsilon}. \tag{7}$$

Substituting into (4) yields

$$y = AK^\alpha \left(\int_\varepsilon \varepsilon^{\frac{1}{1-\alpha}} f(\varepsilon) d\varepsilon \right)^{1-\alpha}. \tag{8}$$

This is a standard aggregate production function, AK^α, augmented by a term that captures a love of variety effect from the optimal allocation of capital across plants. With a given distribution $f(\cdot)$ the idiosyncratic shocks magnify average aggregate productivity.

The condition for **intertemporal optimality** is $u'(c) = \beta E V_K(A', K')$ so that the marginal cost and expected marginal gains of additional capital are equated. Using (1), this condition becomes

$$u'(c) = \beta E u'(c') \left[(1-\delta) + A' \alpha K'^{\alpha-1} \left(\int_\varepsilon \varepsilon'^{\frac{1}{1-\alpha}} f'(\varepsilon) d\varepsilon \right)^{1-\alpha} \right]. \tag{9}$$

The left side is the marginal cost of accumulating an additional unit of capital. The right side is the discounted marginal gain of capital accumulation. Part of this gain comes from having an extra unit of capital to allocate across production sites in the following period. The productivity from these production sites depend on two factors, the **future** values of: aggregate productivity, A' and the cross sectional distribution of idiosyncratic shocks, $f'(\varepsilon)$.

The choice of k for each plant within a period is independent of the choice between consumption and saving. The planner optimally allocates capital to maximize the level of output and then allocates output between consumption and capital accumulation. Clearly, once we allow for limits to reallocation, the capital accumulation decision will depend upon the future allocation of capital across production sites. In this way, variations in the distribution of $f(\cdot)$ can impact on the capital accumulation choice.

2.2 Aggregate Output and Productivity

For this economy, there is an interesting way to represent total output. This is seen from defining

$$\tilde{A} \equiv A \int_{\varepsilon} \varepsilon k(\varepsilon)^{\alpha} f(\varepsilon) d\varepsilon \tag{10}$$

so that

$$y = \tilde{A} K^{\alpha}. \tag{11}$$

from (4). Researchers interested in measuring TFP from the aggregate data will typically uncover $\tilde{A}$ rather than A. This is the mis-measurement referred to earlier. As the discussion progresses, we will refer to $\tilde{A}$ as mis-measured TFP. There are three factors which influence $\tilde{A}$. The first one is A. The influence of A, aggregate TFP, on $\tilde{A}$, measured TFP, is direct and has been central to many studies of aggregate fluctuations. Second, the distribution $f(\varepsilon)$. Variations in $f(\varepsilon)$ influence $\tilde{A}$ because variations in the cross sectional distribution of the idiosyncratic shocks lead to different marginal productivities of plants and thus changes in measured TFP. Finally, there is the allocation of factors, k. If factors are optimally allocated, then the distribution of capital over plants does not have an independent effect on $\tilde{A}$. However, the existence of frictions may imply that, in a static sense, capital is not efficiently allocated. In that case, even with $f(\varepsilon)$ fixed, the reallocation process will lead to variations in $\tilde{A}$.[3]

Starting with Olley and Pakes (1996), many researchers have recognized the dependence of aggregate productivity on factor allocation. In many studies the underlying frictions are

[3]This decomposition of productivity taken from Olley and Pakes (1996) makes clear the interaction between the distribution of productivity and factors of production across firms. Gourio and Miao (2006) use a version of this argument, see their equation (45), to study the effects of dividend taxes on productivity. Khan and Thomas (2006) study individual choice problems and aggregation in the frictionless model with plant specific shocks. Basu and Fernald (1997) also discuss the role of reallocation for productivity in an aggregate model.

due to policies which influence steady state productivity across countries.[4] Our analysis differs from these studies in a couple of important ways. We next focus on (i) frictions through adjustment costs to capital, (ii) dynamic inefficiency brought about through the adjustment process so that the magnitude of the inefficiency and thus aggregate productivity are endogenous and (iii) the behavior of aggregate productivity over business cycles.

3 Capital Adjustment Costs

The allocation of capital over sites with heterogeneous idiosyncratic productivity has important effects on measured total factor productivity. In a frictionless economy there are no cyclical effects of reallocation effects on productivity. However, there is ample evidence in the literature for both non-convex and convex adjustment costs. Introducing these adjustment costs will enrich the analysis of productivity and reallocation.

There are two distinct frictions to study, corresponding to the two dimensions of capital adjustment. The first is "costly reallocation" in which the friction is associated with the allocation of capital across the production sites. The second is "costly accumulation" in which the adjustment cost refers to the cost of accumulating rather than allocating capital.

Our focus here lies on studying the presence of costs to the reallocation (assignment) process. We introduce a special type of adjustment costs that is very tractable, although not very informative about the source of the friction, following Calvo (1983).[5] Each period a Bernoulli draw determines the fraction $\pi \in [0, 1]$ of plants the planner can costlessly reallocate capital between. The remaining fraction of plants $1-\pi$ produces with its beginning-of-period capital stock. This structure of adjustment costs captures the fact that plants adjust their capital stock infrequently.

By assumption, π is not dependent on the state of the plant. This simplification makes our analysis tractable. At the same time it does not preclude a role for the cross sectional distribution in the state space of the problem. Presumably, allowing state dependent adjustment rates could add another form of non-linearity to the model. As the adjustment rate is fixed, some smoothing in response to shocks through adjustment on the extensive margin is precluded.

Applying a law of large numbers, the plant-specific shocks ε are assumed to be equally distributed over the fractions π and $1 - \pi$ of adjustable and non-adjustable plants. The two distributions of plants will be referred to as F^a and F^n. This also implies that $E(\varepsilon)$ is time-invariant and the same across adjustable and non-adjustable plants.

[4] Bartelsman, Haltiwanger, and Scarpetta (2006) discuss these other studies in their analysis of productivity differences over 24 economies.

[5] Sveen and Weinke (2005) adopt a similar structure.

3.1 The Planner's Problem

For the dynamic program of the planner in the presence of adjustment costs, the state vector contains aggregate productivity A, the aggregate capital stock K, and Γ. The high-dimensional object Γ describes the joint distribution over capital (at the start of the period) and productivity shocks across plants. Γ is needed in the state vector because the presence of adjustment costs implies that a plant's capital stock may not reflect the current draw of ε. As noted above, there is time variation in the probability of adjustment π. Furthermore, there are shocks to the variance of idiosyncratic productivity shocks, parameterized by λ. Changes in the variance of the cross-sectional idiosyncratic productivity, as recently highlighted in Bloom (2009) and Gilchrist, Sim, and Zakrajsek (2010), have an effect on output. Such changes can be interpreted as variations in uncertainty. Consider a mean-preserving spread (MPS) in the distribution of ε. In a frictionless economy such a spread would incentivize the planner to carry out more reallocation of capital between plants because capital can be employed in highly-productive sites. Let $s = (A, K, \Gamma, \lambda, \pi)$ denote the vector of aggregate state variables. The adjustment state of a plant is given by $j = a, n$, where a stands for 'adjustment', while n stands for 'non-adjustment'.

Given the state, the planner makes an investment decision K' and chooses how much capital to reallocate across those plants whose capital stock can be costlessly reallocated, $(k, \varepsilon) \in a$. Let $\tilde{k}_j(k, \varepsilon, s)$ for $j = a, n$ denote the capital allocation to a plant that enters the period with capital k and profitability shock ε in group j after reallocation. The capital of a plant in group $j = a$ is adjusted and is optimally set by the planner to the level $\tilde{k}_a(k, \varepsilon, s)$. The capital of a plant in group $j = n$ is not adjusted so that $\tilde{k}_n(k, \varepsilon, s) = k$.

The choice problem of the planner is:

$$V(A, K, \Gamma, \lambda, \pi) = max_{\tilde{k}_a(k,\varepsilon,s),K'} \; u(c) + \beta E_{[A',\Gamma',\lambda',\pi'|A,\Gamma,\lambda,\pi]} V(A', K', \Gamma', \lambda', \pi') \tag{12}$$

subject to the resource constraint (2) and

$$y = \int_{(k,\varepsilon)\in F^a} A\varepsilon\tilde{k}_a(k, \varepsilon, s)^\alpha d\Gamma(k, \varepsilon) + \int_{(k,\varepsilon)\in F^n} A\varepsilon\tilde{k}_n(k, \varepsilon, s)^\alpha d\Gamma(k, \varepsilon), \tag{13}$$

which is simply (4) split into adjustable and non-adjustable plants. Here F^j is the set of plants in group $j = a, n$. The fraction of plants whose capital stock can be adjusted is equal to π

$$\int_{(k,\varepsilon)\in F^a} f(\varepsilon)d\varepsilon = \pi \tag{14}$$

and the amount of capital over all plants must sum to total capital K:

$$\pi \int_{(k,\varepsilon)\in F^a} \tilde{k}_a(k,\varepsilon,s)d\Gamma(k,\varepsilon) + (1-\pi)\int_{(k,\varepsilon)\in F^n} \tilde{k}_n(k,\varepsilon,s)d\Gamma(k,\varepsilon) = K. \tag{15}$$

There is time to build in the model so that new investment increases the capital stock in the following period. As the capital is plant specific, it is necessary to specify transition equations at the plant level. Let $i = \frac{K'-K}{K}$ denote the gross investment rate so that $K' = (1-\delta+i)K$ is the aggregate capital accumulation equation. To distinguish reallocation from aggregate capital accumulation, assume that the capital at **all** plants, regardless of their reallocation status, have the same capital accumulation. The transition for the capital (after reallocation) this period and the initial plant-specific capital next period is given by

$$k'_j(k,\varepsilon,s) = (1-\delta+i)\tilde{k}_j(k,\varepsilon,s), \tag{16}$$

for $j = a, n$. Due to the presence of frictions $\tilde{k}_a(k,\varepsilon,s)$ is not given by (7). Notice that A affects unadjustable and adjustable plants in the same way. This implies that the optimal reallocation decision will occur independently of A. The shock to A will have an effect on the mis-measured part of TFP only in the presence of a capital accumulation problem, since the total amount of capital in adjustable and non-adjustable plants may differ.

The quantitative analysis will focus on reallocation of capital, defined as the fraction of total capital that is moved between adjustable plants within a period. Following a new realization of idiosyncratic productivity shocks, the planner will reallocate capital from less productive to more productive sites. Aggregate output is thus increasing in the amount of capital reallocation.

As $\tilde{k}_a(k,\varepsilon,s)$ denotes the post-reallocation capital stock of a plant with initial capital k, the plant-level reallocation rate would be $r(k,\varepsilon,s) = |\frac{\tilde{k}_a(k,\varepsilon,s)-k}{k}|$. Aggregating over all the plants who adjust, the aggregate reallocation rate is

$$R(s) \equiv 0.5 \int_{(k,\varepsilon)\in F^a} r(k,\varepsilon,s)d\Gamma(k,\varepsilon). \tag{17}$$

The multiplication by 0.5 is simply to avoid double counting flows between adjusting plants.

3.2 Joint Distribution of Capital and Productivity

In the presence of reallocation frictions, the state space of the problem includes the cross sectional distribution, Γ. Consequently, when making investment and reallocation decisions the planner needs to forecast Γ'. It is computationally not feasible to follow the joint distribution of capital and profitability shocks over plants, we represent the joint distribution with several of its moments. These forecast the marginal benefit of investment.

The right set of moments is suggested by the following expression for aggregate output. Using $E(XY) = E(X)\cdot E(Y) + Cov(X,Y)$ to rewrite (13) yields:

$$y = \pi(\bar{\varepsilon}\mu_a + \phi_a) + (1-\pi)(\bar{\varepsilon}\mu_n + \phi_n), \tag{18}$$

where $\mu_j \equiv E(\tilde{k}_j(k,\varepsilon,s)^\alpha)$ and $\phi_j \equiv Cov(\varepsilon, \tilde{k}_j(k,\varepsilon,s)^\alpha)$, for $j = a, n$. Instead of Γ we retain μ_n and ϕ_n in the state vector.

These two moments contain *all* the necessary information about the joint distribution of capital and profitability among non-adjustable plants. The information about capital in plants $\in F^A$ is not needed since capital in those plants can be freely adjusted, independently of their current capital stock. Together, μ_n and ϕ_n are sufficient to compute the output of those plants whose capital cannot be reallocated and thus to solve the planner's optimization problem. Note that by keeping μ_n and ϕ_n in the state space, we are not *approximating* the joint distribution over capital and productivity since the two moments can account for all the variation of the joint distribution. This feature of our choice of moments allows us to compare it with common approximation techniques in the spirit of Krusell and Smith (1998).

The covariance term ϕ_n is crucial for understanding the impact of reallocation on measures of aggregate productivity. If the covariance is indispensable in the state vector of the planner, then the model is not isomorphic to the stochastic growth model. That is, if the covariance is part of the state vector, then the existence of heterogeneous plants along with capital adjustment costs matters for aggregate variables like investment over the business cycle.

When either A or π is stochastic, it is possible to follow the evolution of these moments analytically.[6] The choice of $\tilde{k}_a$ for adjustable plants, along with the respective ε shocks at these plants, maps into values of the moments μ_a and ϕ_a. Together with the new realization of exogenous shocks at the beginning of the next period these map into the next period moments μ_n' and ϕ_n'. The laws of motion for the two states μ_n and ϕ_n are given by

$$\mu_n' = \pi'\mu_a + (1-\pi')\mu_n \tag{19}$$

and

$$\phi_n' = \pi'\rho_\varepsilon\phi_a + (1-\pi')\rho_\varepsilon\phi_n. \tag{20}$$

Together these laws of motion define the law of motion of the joint distribution Γ, allowing us to follow the evolution of this component of the aggregate state.[7] Equations (18)-(20) permit us to study the trade-off regarding the optimal allocation of capital across sites. The planner can increase contemporaneous output by reallocating capital from low- to high-productivity sites in F^a. This will increase the covariance between profitability and capital, ϕ_a, while at

[6] The analytics hold for the evolution of the mean, (19), but not the covariance, (20), when λ is stochastic.

[7] Note that $\phi' = Cov(k(\varepsilon)^\alpha, \varepsilon')$ is an expectation. The term ε' is made up of two components, one is the persistent part, and one is an i.i.d. part, denoted η. Rewrite $\varepsilon' = \rho_\varepsilon\varepsilon + (1-\rho_\varepsilon)\eta$ to obtain $\phi' = Cov(k(\varepsilon)^\alpha, \rho_\varepsilon\varepsilon + (1-\rho_\varepsilon)\eta) = \rho_\varepsilon\phi$.

the same time decreasing μ_a because $\alpha < 1$. A fraction $1 - \rho_\pi$ of currently adjustable plants will not be able to adjust its capital stock tomorrow. The planner therefore has to trade off the higher instantaneous output from reallocation with the higher probability of a mismatch between $\tilde{k}_n(k, \varepsilon, s) = k$ and the realization of ε' for plants in F^n tomorrow. This is what is captured in the laws of motion (19) and (20).

3.3 Stationary Equilibria

To fix ideas we can analyze the stationary economy where π and λ are not varying over time. In this environment a stationary distribution Γ^* exists. Using (19) it follows that $\mu_n = \mu_a = \mu^*$. Furthermore, stationary values ϕ_a^* and ϕ_n^* exist. Using (20) one can show that ϕ_n converges to

$$\phi_n^* = \phi_a^* \frac{\pi \rho_\varepsilon}{1 - (1 - \pi)\rho_\varepsilon}. \tag{21}$$

Hence (18) becomes

$$y = \bar{\varepsilon}\mu^* + \Lambda \phi_a^*, \tag{22}$$

where $\Lambda \equiv \frac{\pi}{1-(1-\pi)\rho_\varepsilon}$ is a function of parameters. Λ is (weakly) increasing in both π and ρ_ε.[8] Intuitively, an increase in π increases total output because more plants' capital stock can be costlessly adjusted. An increase in ρ_ε, the persistence of idiosyncratic productivity shocks, implies that the probability of a plant switching status and being non-adjustable with a mismatch between ε and k is decreased.[9]

Figure 1 shows equilibrium values of μ^* and ϕ_a^* in stationary economies for different values of π. As $\pi \to 0$ the planner reallocates less capital between plants. A value of $\mu^* = 1$ implies $\phi_a^* = 0$, because $k(\varepsilon) = 1$ for all sites, meaning that the capital level is independent of ε. On the other hand, as the fraction of adjustable plants increases, the optimal ϕ_a^* increases.

4 Quantitative Results

With exogenous movements in π and λ no stationary distribution of Γ exists and the two moments μ_n and ϕ_n become part of the state vector. This problem can no longer be solved analytically. This section presents quantitative results.

In the stationary economy, reallocation effects only mattered for aggregate **levels**. When are reallocation effects likely to play a role for aggregate **dynamics**? One key prerequisite

[8]Formally, $\frac{\partial \Lambda}{\partial \pi} = \frac{1-\rho_\varepsilon}{[1-(1-\pi)\rho_\varepsilon]^2} \geq 0$, $\frac{\partial \Lambda}{\partial \rho_\varepsilon} = \frac{\pi(1-\pi)}{[1-(1-\pi)\rho_\varepsilon]^2} \geq 0$. The cross-derivatives are given by $\frac{\partial^2 \Lambda}{\partial \rho_\varepsilon \partial \pi} = \frac{\partial^2 \Lambda}{\partial \pi \partial \rho_\varepsilon} = \frac{1}{[1-(1-\pi)\rho_\varepsilon]^2} - \frac{2\pi}{[1-(1-\pi)\rho_\varepsilon]^3}$.

[9]In the extreme case of iid shocks to idiosyncratic productivity shocks the planner would be more reluctant to allocate large amounts of capital to high-productivity sites, decreasing aggregate output.

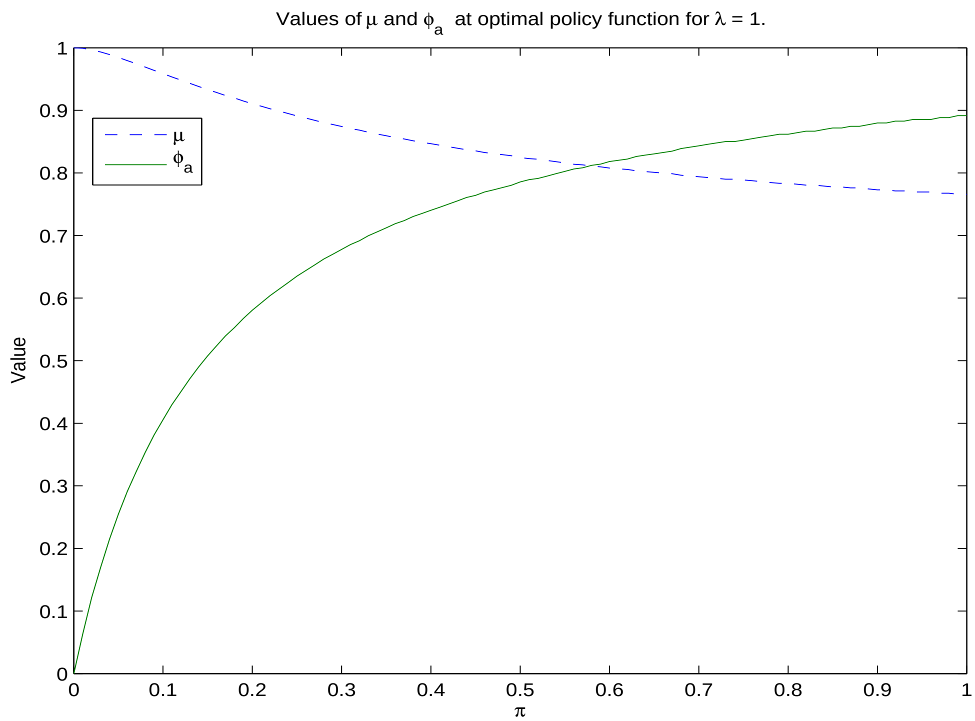

Figure 1: Values of μ and ϕ_a in stationary equilibrium for various π. Economy with $\lambda = 1$ and $\rho_\varepsilon = .9$

is that the economy be subject to shocks that cause the distribution Γ to move over time. Without movements in Γ the benefits from reallocation are constant and the covariance term ϕ is not required to forecast Γ'. The reasons why Γ may vary and the implications of its variability will be clear as the analysis proceeds.

In keeping with the distinction noted earlier between reallocation and accumulation, the initial quantitative analysis, presented in section 4.1 is for an economy with a fixed capital stock, thus highlighting reallocation. The economy is then enriched to allow for capital accumulation in section 4.2.

For each of these models, this section focuses on the effects of capital reallocation on aggregate productivity. In addition, we present evidence on whether higher order moments are needed in the solution of the planner's optimization problem in Section 5. As highlighted in the introduction, these two themes are connected: higher order moments are needed to follow the evolution of Γ precisely when capital reallocation matters for the cyclical movements in productivity.

We solve the model at a quarterly frequency, using these **baseline** parameters. We model an economy with $N = 10,000$ plants and a labor share of $\alpha = 0.35$. We assume log-utility. Assuming an annual interest rate of 4% this implies a discount factor $\beta = 0.987$ and a depreciation rate $\delta = 0.025$. We set the mean of π to $\bar{\pi} = 0.5$. This implies that plants adjust their capital stock on average every two quarters. Sveen and Weinke (2005) treat changes in the capital stock of under 10% in absolute value as maintenance and hence use $\pi = 0.08$. In our setup, the choice of π mainly affects aggregate *levels*, not transitions. Aggregate profitability takes the form of an AR(1) in logs

$$\ln a_t = \rho_a \ln a_{t-1} + \nu_{a,t}, \qquad \nu_a \sim N(0, \sigma_a), \tag{23}$$

where $\rho_a = 0.9$ and $\sigma_a = 0.005$. Idiosyncratic profitability shocks are log-normal and evolve according to a law of motion with time-varying variance

$$\ln \varepsilon_t = \rho_\varepsilon \ln \varepsilon_{t-1} + \lambda_t \nu_{\varepsilon,t}, \qquad \nu_{\varepsilon,t} \sim N(0, \sigma_\varepsilon). \tag{24}$$

The parameters of the idiosyncratic shock process are $\rho_\varepsilon = 0.9$ and $\sigma_\varepsilon = 0.2$. The parameter λ governs the mean-preserving spread of the normal distribution from which idiosyncratic profitability ε is drawn. It has a mean of 1 and variance σ_λ

$$\lambda_t = \rho_\lambda \lambda_{t-1} + \nu_{\lambda,t}, \qquad \nu_{\lambda,t} \sim N(1, \sigma_\lambda) \tag{25}$$

Finally, the process of π follows

$$\pi_t = \rho_\pi \pi_{t-1} + \nu_{\pi,t}, \qquad \nu_{\pi,t} \sim N(\bar{\pi}, \sigma_\pi). \tag{26}$$

We set $\rho_\lambda = 0.82$ as in Gilchrist, Sim, and Zakrajsek (2010) and $\rho_\pi = 0.9$. In order to be able to compare the effect of different shocks, the standard deviation of the innovations, $\sigma_\pi = 0.03$ and $\sigma_\lambda = 0.014$ are set to generate the same amount of variation in output as shocks to a.

Section 4.3 explores the sensitivity of our findings to this parameterization. The computational strategy is discussed in the Appendix.

4.1 Capital Reallocation

Table 1 shows measures of the efficiency of the allocation of capital and the cyclicality of mis-measured TFP. These two aspects of the economy are inherently linked. Aggregate productivity is endogenous and responds to changes in the amount of capital reallocated.

The column labeled 'R' for 'Reallocation' measures the time series average of the cross-sectional reallocation of capital across plants as defined in (17), relative to the frictionless benchmark without exogenous shocks. The column labeled $E_t(\sigma_i(arpk_{it}))$ measures the time series average of the cross sectional standard deviation of the average marginal product of capital. Column σ_y/μ_y shows the coefficient of variation of output. The column labeled G shows the output gap, defined as $G(s) = \frac{y^{FL}(s) - y(s)}{y^{FL}(s)}$. The column reports output in state s

relative to the frictionless benchmark.[10] The column labeled $\sigma(\tilde{A}/A)$ reports the standard deviation of measured relative to actual TFP. The last column $c(R, \tilde{A})$ shows the correlation between the time series for capital reallocation, R, and mis-measured productivity, $\tilde{A}$.

The first block of Table 1 reports results for the frictionless economy. The second block of results introduces capital adjustment costs.

Case	R/R^*	$E_t(\sigma_i(arpk_{it}))$	σ_y/μ_y	G	$\sigma(\tilde{A}/A)$	$C(R, \tilde{A})$
Frictionless						
nonstochastic	1 (0)	0 (−)	0 (−)	0 (−)	0 (−)	*na* (−)
stochastic A	1 (0)	0 (−)	0.011 (0.0007)	0 (−)	0 (−)	*na* (−)
stochastic λ	1 (0)	0 (−)	0.011 (0.00016)	0 (−)	0.024 (0.0004)	0.955 (0.004)
Frictions						
nonstochastic	0.477 (0)	0.659 (0)	0 (−)	0.045 (0)	0 (−)	*na* (−)
stochastic A	0.477 (0)	0.659 (0)	0.011 (0.001)	0.045 ($1.3e{-}09$)	0 (−)	*na* (−)
stochastic π	0.476 (0.004)	0.661 (0.002)	0.011 (0.0003)	0.047 (0.0003)	0.021 (0.0006)	0.993 (0.003)
stochastic λ	0.477 (0.0001)	0.660 (0.002)	0.009 (0.0002)	0.046 ($2.8751e{-}05$)	0.018 (0.0004)	0.882 (0.007)
stochastic π, λ	0.478 (0.003)	0.654 (0.005)	0.015 (0.0003)	0.046 (0.0004)	0.030 (0.0007)	0.848 (0.008)

Table 1: Capital Reallocation Model: Productivity Implications

Results from 100 simulations with T=1000, standard deviations in parentheses below. $\frac{R}{R^*}$ measures the time series average of the cross-sectional reallocation of capital across plants, relative to the frictionless benchmark, R^*. $E_t(\sigma_i(arpk_{it}))$ is the mean standard deviation of the average revenue product of capital. σ_y/μ_y is the coefficient of variation of output. G refers to the output gap relative to the frictionless benchmark. The column $\sigma(\tilde{A}/A)$ shows the standard deviation of measured vs. real TFP. The last column $c(R, \tilde{A})$ shows the correlation between capital reallocation and mismeasured TFP. The "*na*" entry means that the correlation is not meaningful as one of the variables is constant.

4.1.1 Frictionless Economy

The first row of Table 1 shows the results for the frictionless economy, $\pi = 1$, without time series variations in TFP, the volatility of the idiosyncratic shocks λ, or the fraction of adjustable sites π.[11] This case serves as a benchmark. Without frictions, the marginal

[10] The frictionless output $y^{FL}(s)$ is a function of s because changes in λ affect the output achieved in the frictionless case.

[11] In this abbreviated problem, the planner solves $V(\Gamma) = max_{k(\varepsilon)}\ u(c) + \ \beta EV(\Gamma')$ subject to the resource constraint (2) and total production given by (13).

product of capital is equalized across plants and our measure of the inefficiency of the capital allocation, $E_t(\sigma_i(arpk_{it}))$, is zero. The first-best output is achieved. The mis-measurement of TFP is constant. The amount of capital reallocation is time-invariant and hence plays no role for aggregate productivity.

The second row, 'stochastic A' introduces variation in aggregate profitability. Variations in A have no effect on the reallocation of capital in this economy, because the planner reallocates capital across plants within a period. Consequently the amount of reallocation is the same as without variations in A. The allocation is efficient, $\tilde{A}$ varies only with A and the output gap is zero. The only difference with respect to the benchmark in the previous row is the variability of output, which is affected by changes in aggregate profitability. Since A enters total output multiplicatively $\sigma_y/\mu_y = \sigma_a/\mu_a$, i.e. all variation in output stems from variation in A. As before, the amount of capital reallocation is time-invariant.

The third row 'stochastic λ' presents results for the frictionless economy with stochastic variance of idiosyncratic productivity shocks. The parameter λ is chosen to generate the same coefficient of variation of output as the previous case.[12] The resulting allocation is always efficient, as reflected in the output gap of zero. Importantly, output and mis-measured TFP vary with λ, as shown in columns σ_y/μ_y and $\sigma(\tilde{A}/A)$. This represents a pure reallocation effect through changes in $f(\varepsilon)$ and occurs even under constant A and π. The last column shows the high correlation between the amount of capital reallocation and output. In periods where more capital is reallocated between adjustable plants, aggregate output is higher. The correlation is less than 1 because of the persistence of shocks to λ. In periods where λ changes, intra-period reallocation reacts, but output is still affected by the current allocation of capital across non-adjustable plants.

This economy presents the simplest case where reallocation is the sole driver of business cycles. To some degree, it looks like an economy driven by exogenous TFP. Here the variations in productivity arise from the endogenous reallocation of capital. The following subsection studies environments where capital adjustment costs amplify this feature.

4.1.2 Costly Capital Reallocation

Setting $\pi < 1$ introduces capital adjustment costs to the frictionless economy, so that only a fraction of all plants' capital stocks can be adjusted within a given period. Costly capital reallocation will have effects on measured productivity and its cyclical properties.

When π is non-stochastic and there are no other aggregate shocks, a stationary joint distribution Γ exists, with the moments (μ_n, ϕ_n) constant, as was shown in Section 3.3 above.

[12]For this case, λ takes values between 0.966 and 1.0344. These values are chosen to generate the same amount of output volatility as direct shocks to a. Below we study the implications of larger variability in λ. Note that $\lambda > 1$ can imply that some values of the shock become negative. To avoid this, we apply the MPS to the underlying normal distribution and re-adjust its mean such that mean of the log-normal is preserved.

Table 1 shows the results for this case in the row labeled 'nonstochastic'. In this economy the fraction of capital reallocated is far below the frictionless benchmark, as indicated in the second column. With $R < \pi$, the planner's chosen distribution of capital over adjustable plants is different from the distribution in the frictionless case. Although capital in a fraction π of plants could be costlessly reallocated, the reallocation rate is less than π. Instead, reallocation is lower indicating a reduced capital flow beyond the direct influence of $\pi < 1$.

Figure 2 plots capital reallocation as a function of π. The dashed line is the 45° line. The concave solid green line above it shows capital reallocation between adjustable plants (as a fraction of the frictionless benchmark). As $\pi \to 1$ it approaches the allocation derived in (7). For total capital reallocation (plotted as the red solid line beneath the 45° line) this implies that it approaches π as $\pi \to 1$.

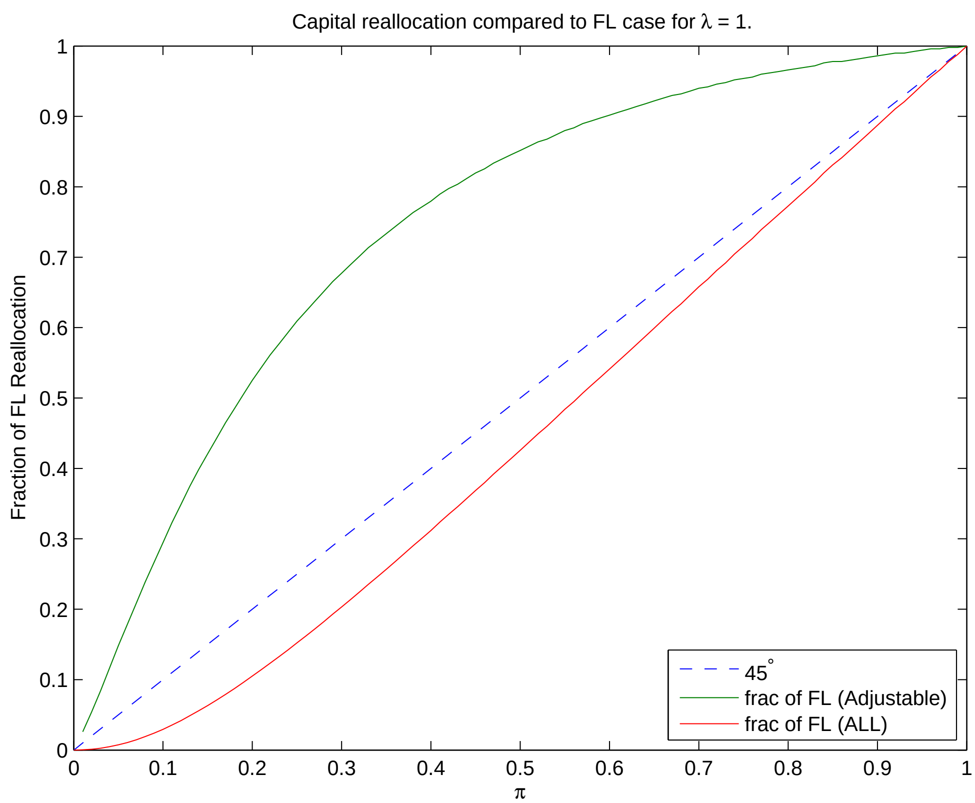

Figure 2: Capital Reallocation in adjustable and all plants as fraction of frictionless benchmark in stationary equilibrium for various π. Economy with $\lambda = 1$ and $\rho = .9$.

The inefficiency of the allocation when $\pi < 1$ is highlighted by the column labeled $E_t(\sigma_i(arpk_{it}))$. This measure of the inefficiency of the allocation is larger than zero, reflecting frictions in the reallocation process that stem from two sources. First, the planner chooses not to equalize marginal products between adjustable plants, reflecting the tradeoffs discussed above. Secondly, the marginal products of capital among non-adjustable plants exhibit a high degree of heterogeneity due to the fact that their capital is fixed despite a new realization of idiosyncratic profitability. Column four shows that because ϕ_n and μ_n converge to their steady-state values output does not vary in this economy. The output gap is positive now, directly reflecting the impact of $\pi < 1$. Importantly, the mis-measurement in TFP is constant over time, we only obtain a level-effect.

The row labeled 'stochastic A' allows for randomness in aggregate productivity with constant π. As explained above, the amount of reallocation is independent of variations in A. Output and $\tilde{A}$ vary only with A. Because $\pi < 1$ the allocation is characterized by a positive standard deviation of marginal and average revenue products of capital and a positive output gap.

Variations in π create time series variation in the moments μ_n and ϕ_n, as shown in the row 'stochastic π'. This economy generates the same coefficient of variation of output as the cases discussed above. Fluctuations in π lead to pro-cyclical capital reallocation patterns. From the last column the correlation between reallocation and output is very high. But this is not *simply* a correlation. In the presence of adjustment frictions, reallocation **causes** the observed time-variations in output. Variations in π therefore also lead to variations in (mis-measured) total factor productivity. There is considerable misallocation as a result of $\pi < 1$.

The marginal products of capital are not equalized across plants, neither among the adjustable nor the unadjustable sites. This results in a positive output gap which varies with the evolution of μ_n and ϕ_n. This gap is about 4.5% of real GDP. Additionally, this economy exhibits counter-cyclical productivity dispersion. When π is low, less capital can be reallocated between adjustable plants. This decreases output and increases the standard deviation of marginal products between those plants. Though λ is held fixed, $\sigma_i(arpk_{it})$ nonetheless varies over time.

Figure 3 shows an impulse responses for a negative shock to π in period $t = 5$. The x-axes show time, while the y-axes in Panels 2-4 shows the % deviation from the unconditional mean. The drop in π is plotted in the first panel, while λ and A are set to their unconditional means. The second panel shows the evolution of the two moments μ_n and ϕ_n. The negative correlation between the two series is very high, as changes in π effect the evolution of μ_n and ϕ_n in very similar ways. The third panel illustrates the co-movement between reallocation 'R' and mis-measured TFP, while the last panel shows investment and output.

The row 'stochastic λ' of Table 1 studies the effects of time-variation in $f(\varepsilon)$ under costly capital reallocation. Due to the presence of adjustment costs, the marginal products of capital cannot be equalized over time. In addition, the variations in λ lead to changes in the optimal allocation decision by the planner and create considerable time-variation in μ_n and ϕ_n. The resulting fluctuations in output are the outcome of different reallocation choices

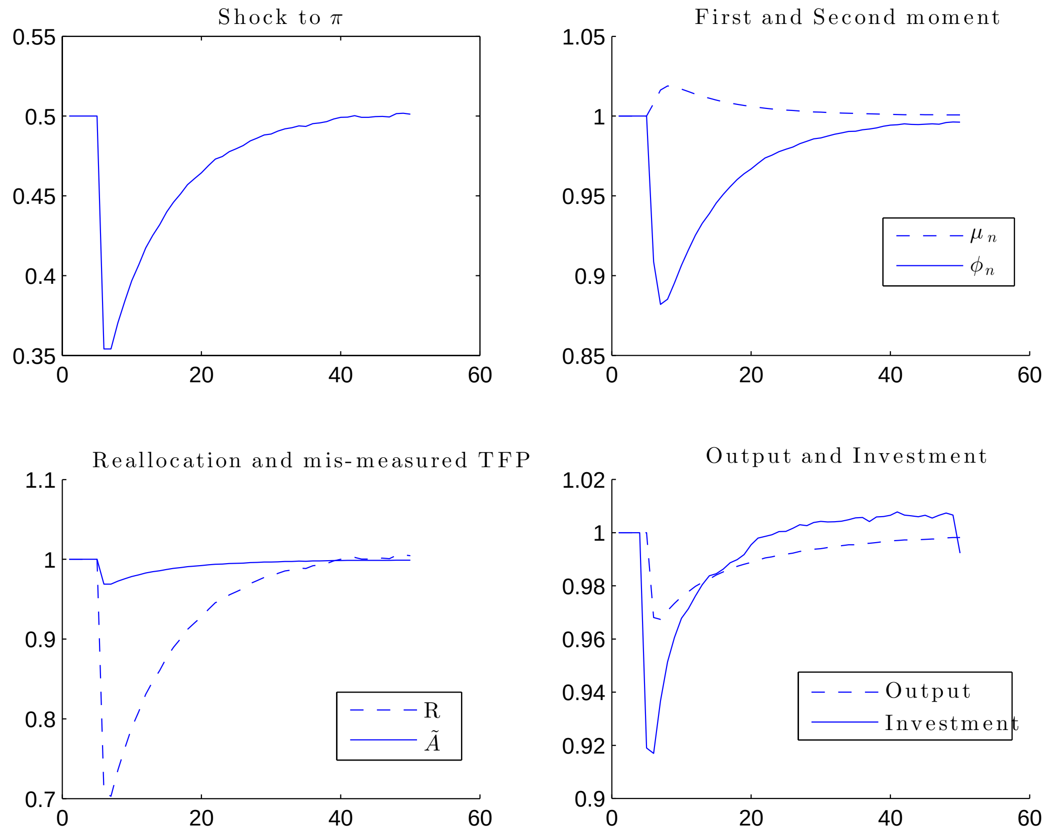

Figure 3: Variations in π: Impulse Response Functions. The y-axes show % deviations from unconditional means.

of the planner that reflect in variations of mis-measured TFP. Compared to the frictionless economy with stochastic λ (row 3) the the model now produces a sizable output gap.

While variations in π affect output directly through the fraction of plants among which capital can be reallocated, the effect of changes in λ is less direct. Variations in λ induce different reallocation choices but a fraction of the effect on output comes from the fact that the marginal revenue product of capital is changed through productivity draws with larger or smaller tails.

Figure 4 shows an impulse responses for a negative shock to λ. The reversion of λ to its unconditional mean after the shock is plotted in the first panel, while π and A are set to their unconditional means. The second panel shows the evolution of the two moments μ_n and ϕ_n. The sharp drop in ϕ_n is a direct effect of the shock to λ, whereas the increase in μ_n reflects the effects of different reallocation choices. The third panel illustrates the co-movement between reallocation 'R' and mis-measured TFP, while the last panel shows investment and output.

Panel 2 shows the effect of variations in λ on the two moments of the cross sectional distribution, μ_n and ϕ_n. The panel highlights that Γ is a slow moving state variable, implying that μ_n and ϕ_n do not adjust immediately to their new values following a change in λ. Furthermore, the variations in λ have different effects on ϕ_n (direct) and μ_n (indirect), making the two moments imperfectly correlated. Variations in λ produce more cyclicality in ϕ_n than in μ_n, as was conjectured.

Panel 3 shows the connection between mis-measured TFP and reallocation, which leads to a cyclical effect on output. In this economy with time-varying idiosyncratic uncertainty in the presence of adjustment costs there is a strong cyclical dimension of capital reallocation. Reallocation is driving time-variations in output. For this simulation the correlation between mis-measured TFP and reallocation was 0.977.

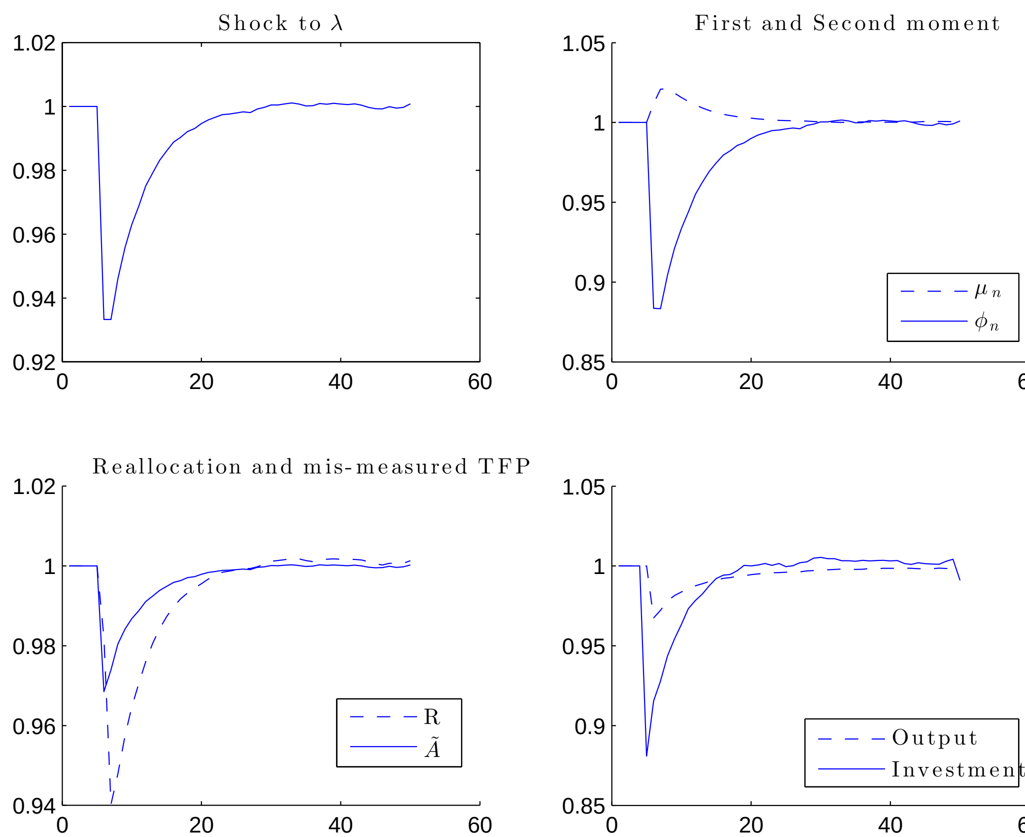

Figure 4: Variations in λ: Impulse Response Functions. The y-axes show % deviations from unconditional means.

The joint effects of changes in π and λ are presented in the last row of Table 1. Output varies significantly over time, with variations resulting directly from both shocks to π and λ. While $\pi < 1$ leads to a positive output gap the presence of a stochastic λ causes additional variation in this gap as was the case before. Notably, mis-measured TFP exhibits significantly more time variation than in the cases of varying λ or varying π alone. This is the result of changes in π and λ jointly affecting the slow-moving joint distribution Γ. Importantly, the correlation between capital reallocation and output is much lower in this environment. This comes about because mis-measured TFP reacts more strongly through changes in λ than π. On the other hand, both exogenous shocks affect the amount of reallocation. The effect of varying π on reallocation, however, is predominantly an extensive margin effect, as a changing fraction of plants can reallocate capital. The effect of λ is an intensive margin effect: more capital is reallocated within a given fraction of adjustable plants. Together this explains the observed decrease in the correlation between reallocation and output.

Overall, adjustment frictions reduce reallocation, generating a non-degenerate distribution of average (and marginal) products of capital across plants. The cost is a reduction in output of about 4.5%, relative to the frictionless benchmark. In all of the experiments, reallocation is procyclical. For these cases, measured variations in TFP are the consequence of reallocation rather than true variations in aggregate productivity. Variations in π lead to countercyclical productivity dispersion across firms.

4.2 Endogenous Capital Accumulation

With endogenous capital accumulation, solving (12), the capital *reallocation* process has significant interactions with the capital *accumulation* decision. The frictions exert a level effect on the optimal capital stock and induce different dynamics following an exogenous shock. As we saw above, reallocation behaves cyclically in the presence of time-series variation in π and/or λ. Variations in λ and π affect the instantaneous value of existing capital and, because of persistence, the expected future return to capital, too. This affects the planner's incentives to invest. Even absent any frictions to capital accumulation the dynamics of investment and consumption are considerably altered by the presence of exogenous shocks to reallocation or the variance of the idiosyncratic shock.

Adding endogenous capital accumulation does not alter the results on the reallocation process shown in Table 1. The reason parallels the argument for the independence of reallocation from A. From (8), total output is proportional to AK^{α}. Thus just as variations in A scale moments, so will variations in K. Consequently, the analysis focuses on the effects of frictions in reallocation on capital accumulation.

Table 2 summarizes results for the endogenous capital accumulation problem, using the baseline parameters, defined earlier. The aggregate capital stock is now endogenous and creates additional variation. The average capital stock is shown in the $\bar{K}$ column. The other columns report correlations of reallocation with investment and output, $C(R, i)$ and $C(R, y)$ and the correlation of investment and mis-measured TFP, $C(\tilde{A}, i)$.

Case	$\bar{K}$	$C(R,i)$	$C(R,y)$	$C(\tilde{A},i)$
Frictionless				
nonstochastic	94.4 (0)	*na* (−)	*na* (−)	*na* (−)
stochastic A	94.4 (1.18)	*na* (−)	*na* (−)	*na* (−)
stochastic λ	94.4 (1.05)	0.769 (−)	0.926 (−)	0.8584 (−)
Frictions				
nonstochastic	87.86 (−)	*na* (−)	*na* (−)	*na* (−)
stochastic A	88.01 (1.15)	*na* (−)	*na* (−)	0.752 (−)
stochastic π	87.72 (0.93)	0.709 (−)	0.963 (−)	0.714 (−)
stochastic λ	87.85 (0.73)	0.700 (−)	0.931 (−)	0.779 (−)
stochastic π, λ	87.47 (1.42)	0.734 (−)	0.809 (−)	0.884 (−)

Table 2: Endogenous Capital Accumulation: Aggregate Moments

Results from 100 simulations with T=1000, N=10,000 are reported with standard deviations in parentheses below. Simulations with frictions were computed with a mean of π equal to 0.5, mean of $\lambda = 1$, a ρ of 0.6, N=10,000 plants. $\bar{K}$ reports the average capital stock. $C(R,i)$ is the correlation between reallocation and investment, $C(R,y)$ is the correlation between reallocation and output, and $C(\tilde{A},i)$ is the correlation between mis-measured TFP and investment. The "*na*" entry means that the correlation is not meaningful as one of the variables is constant.

From Table 2, the interaction of costly reallocation and accumulation is evident in a number of forms. First, $\bar{K}$, which is the average capital for a particular treatment, depends on the nature and magnitude of the capital adjustment costs. Even in the absence of any aggregate shocks, the capital stock is almost 9% lower when there are adjustment frictions compared to the frictionless case. This comparison of the average capital stocks with and without frictions stands regardless of the source of the shocks.

Second, the addition of the shocks increases the variability of capital, as indicated by the standard deviation entry for capital. For example, with frictions the standard deviation of the capital stock is 1.42 when both π and λ are random. This is about 25% more than the variability of capital when there are only exogenous productivity shocks.

Third, capital accumulation is positively correlated with both reallocation and mis-measured TFP. An increase in λ, for example, leads to an increase in investment, reallocation and output. The correlation of reallocation and investment, $C(R,i)$, is informative about the effects of frictions on the incentive to accumulate capital.[13] This correlation is highest

[13]For the nonstochastic and stochastic A models, this correlation is not defined as capital reallocation is

when π and λ are random. In this case, an increase in π say, will imply that more plants are able to adjust and for this reason alone reallocation will increase. With π correlated, it is likely that more plants will be able to adjust in the future, so investment increases too. The magnitude of this correlation is smaller when λ is random. Though the same fraction of plants adjusts each period, the gains to adjustment are larger when λ is high. This generates a positive correlation between reallocation and investment.

Finally, reallocation is procyclical in the presence of shocks to either π or λ. This returns to one of the themes of the paper. If variations arise from either changes in the fractions of adjusting plants, through π, or by a change in the spread of the shocks, through λ, output responds. The key to this response is reallocation: the effects on output of getting the right amount of capital into its most productive use. This is captured through $\tilde{A}$.

4.3 Robustness

The previous results illustrated a couple of themes. First, variations in either π or λ are necessary to generate cyclical movements in reallocation, with resulting effects on mis-measured TFP. Second, evolution of the cross sectional distribution generated dynamics only in the stochastic π and/or λ cases. This is illustrated by the fact that higher order moments are relevant in the planner's optimization problem and the evolution of this moments are seen in the impulse response functions.

Parameter changes	R/R^*	$E_t(\sigma_i(arpk_{it}))$	σ_y/μ_y	G	$\sigma(\tilde{A}/A)$	$C(R,\tilde{A})$	$C(R,i)$
Frictions							
Baseline	0.478 (0.003)	0.654 (0.005)	0.015 (0.0003)	0.046 (0.0004)	0.030 (0.0007)	0.848 (0.008)	0.734 (−)
$\bar{\pi} = 0.3$	0.265 (0.005)	0.815 (0.004)	0.024 (0.0008)	0.09 (0.0004)	0.047 (0.002)	0.95 (0.004)	0.829 (−)
$\bar{\pi} = 0.9$	0.897 (0.004)	0.170 (0.006)	0.012 (0.0002)	0.006 (0.0001)	0.025 (0.0005)	0.56 (0.03)	0.488 (−)
$\rho_\pi = 0.5$	0.477 (0.001)	0.657 (0.001)	0.011 (0.0001)	0.046 ($9.5e{-}05$)	0.02 (0.0003)	0.68 (0.008)	0.630 (−)
$\rho_\varepsilon = 0.5$	0.367 (0.008)	1.135 (0.006)	0.020 (0.0005)	0.11 (0.0005)	0.033 (0.0007)	0.97 (0.001)	0.874 (−)
$\sigma_\lambda = 0.1$	0.480 (0.005)	0.696 (0.013)	0.080 (0.003)	0.072 (0.001)	0.136 (0.003)	0.694 (0.02)	0.686 (0.019)
$\sigma_\lambda = 0.1$, $\rho_\lambda = 0.5$	0.480 (0.006)	0.689 (0.006)	0.050 (0.0007)	0.075 (0.002)	0.089 (0.005)	0.524 (0.014)	0.525 (0.012)

Table 3: Capital Reallocation: Robustness

Model with stochastic π and λ. Standard deviations in parentheses.

This section studies the robustness of these findings to alternative values of key param-

constant.

eters. Table 3 reports our findings. It has the same structure as Table 1. The first column indicates the model. The baseline is the case with adjustment costs and stochastic (π, λ) taken from Table 1.

The baseline model assumes $\bar{\pi} = 0.5$. The second and third rows of Table 3 study the implications of lower and higher adjustment rates. Not surprisingly, the reallocation rate is increasing in π, as frictions are lower. This is consistent with Figure 2. The correlation of reallocation and mis-measured TFP is positive, though lower than in the baseline at $\pi = 0.90$.

The standard deviation of actual to mismeasured TFP also varies with $\bar{\pi}$. When $\bar{\pi}$ is high, the response of the planner to a variation in λ is to reallocate capital so that $\sigma(\tilde{A}/A)$ is small compared to the case of low $\bar{\pi}$. This is reflected in the mean standard deviation of the average revenue product of capital.

The table includes two rows in which the serial correlation of shocks is set to 0.5, lower than their baseline values of $\rho_\pi = 0.9$ and $\rho_\varepsilon = 0.9$. Relative to the baseline, the reduction in the serial correlation of π leads to a reduction in the cyclicality of reallocation. With adjustment opportunities less correlated, the costs of reallocating resources that are subsequently mismatched with productivity is higher. Hence reallocation is less correlated with $\tilde{A}$. Also, the correlation of reallocation and investment is lower than in the baseline reflecting the costs of accumulating capital when future adjustment costs are less certain.

When ρ_ε is decreased, the planner has fewer incentives to reallocate capital among adjustable plants. Consequently, the amount of capital reallocation falls (column 1) and the inefficiency of the solution becomes more pronounced. This can be seen in the larger standard deviation of the marginal products of capital and in a higher output gap.

The row labeled $\sigma_\lambda = 0.1$ increases the variability of λ relative to the baseline where $\sigma_\lambda = 0.014$. This spread is closer to that in Bloom (2009) and Gilchrist, Sim, and Zakrajsek (2010). Not surprisingly, this extra volatility in the spread of idiosyncratic shocks leads to much more volatility in $\tilde{A}$ and more variability in output, σ_y/μ_y, relative to the baseline. Reallocation remains procyclical and positively correlated with investment though less compared to the baseline.

The last row shows how a reduction in the serial correlation of λ influences these moments. With a lower serial correlation of the shocks to λ, the correlation between reallocation and investment, though still positive, is considerably lower than the baseline. With less persistent shocks, investment is less responsive to variations in λ and π that provide the motivation for reallocation.

5 Approximation

The previous section showed that the covariance ϕ *matters* for determining the optimal capital allocation. The problem in (12) includes Γ, the joint distribution of (k, ε). Using the first two moments of this distribution, μ_n and ϕ_n, the evolution of Γ can be tracked perfectly. This is important for the planner, who has to forecast the expected future output from non-adjustable plants, $y^{NA'}$. Variations in π and λ generate movements in Γ and hence in y^{NA}. Capital reallocation was tightly linked to changes in the mis-measurement of TFP when stochastic shocks are present.

Movements in Γ may not be captured well by the first moment μ_n alone. Following certain exogenous shocks, the laws of motion for μ_n and ϕ_n can imply very different transition paths for the two moments. While in the frictionless case the two moments were perfectly correlated, this perfect correlation is broken by the existence of time-variation in the adjustment probability π and/or λ. The significance of reallocation effects is related to the forecasting power of ϕ_n.

Relative to the literature starting with Krusell and Smith (1998), this is an important finding. **In particular, this result is distinguished from preceding papers in that for our environment the approximation of the cross sectional distribution requires higher order moments.**

This section makes two points. First, it emphasizes the importance of including the higher order moments in the state vector. From this we can determine how well the evolution of Γ be captured by different subsets of its moments under different cases of stochastic π and λ.

Second, we compare the aggregate outcome of the model against a standard stochastic growth model. This allows us to determine to what extent the reallocation effects influence cyclical properties of the model.

5.1 Goodness of Fit

Table 4 evaluates the importance of the higher order moments.[14] To understand this table, let "DGP" refer to a data set (and moments) created by solving the baseline model (with stochastic π and λ) using (μ, ϕ) in solving the planner's problem. In (12), the planner forecasts y'_n, the output from non-adjustable plants next period. The correctly specified regression model including both moments is given by

$$y_{n,t}^{DGP} = \beta_0 + \beta_1 \mu_{n,t} + \beta_2 \phi_{n,t} + \beta_3 s_t + \varepsilon_t, \tag{27}$$

[14] Only the stochastic model with frictions is explored. The case of "stochastic A" is not of interest as the higher order moments did not matter. For these experiments, the shocks are held fixed to isolate the effects of the approximation.

where s_t includes π_t and λ_t. Estimation results in $\hat{\beta}_0 = 0$, $\hat{\beta}_1 = 1.6487 = \bar{\varepsilon}$, $\hat{\beta}_2 = 1$, and $\hat{\beta}_3 = 0$ with an $R^2 = 1$. The maximum forecast error (MCFE) is zero. As discussed in Den Haan (2010) a problem of R^2 measures to assess the approximation is that observations generated using the true law of motion are used as the explanatory variable. We construct a series $\hat{\hat{y}}_n$ which is using only the approximate law of motion. The forecast error is defined as $\hat{\hat{\varepsilon}}_{t+1} = |\hat{\hat{y}}_{n,t+1} - y_{n,t+1}|$, and the MCFE is the maximum of this series.

Below we study three cases (experiments). The first takes output of the non-adjusting plants from the DGP and regresses it on an intercept, the exogenous state, and the first moment only. Thus this exercise is about approximating the nonlinear solution with a linear representation. The regression model for the linear approximation is given by (27) where we force $\beta_2 = 0$. From Table 4, the linear representation is very accurate if only π is stochastic. When λ is random, the resulting movements in the distribution of shocks leads to much greater significance of the cross sectional distribution in forecasting (decisions do not change in this experiment).

The second case actually solves the planner's problem under the (false) assumption that the model is linear. The resulting decision rules and expectations are model consistent by construction, but not data consistent.[15] The goodness of fit measure is computed from a regression of the output of the non-adjusting plants in the DGP using the model consistent estimators from the linearized approximation. As before, the linear beliefs in the stochastic π case are approximately consistent with the outcome. Again this is not the case when λ is random. For this experiment, the linear forecast rule leads to very different allocative decisions by the planner. Consequently, the R^2 is quite low – movement in the cross sectional distribution are very important.

In the third case, the planner uses the DGP to obtain a linear approximation of the law of motion. With this representation, the planner solves the optimization problem. In this case, the expectations about the evolution of the state vector is consistent with the data, but not with the model. Here, none of the experiments generate a good fit. The planner is simply unable to capture the nonlinear movements in the economy with a linear approximation of the law of motion.

5.2 Comparison to the RBC Model

One of the key findings of Thomas (2002) and the literature that followed was the near equivalence between the **aggregate moments** of a model with lumpy investment and the aggregate implications of a real business cycle model with quadratic adjustment costs at the plant-level. This sub-section returns to that theme of approximating the solution to our

[15] The R^2 from the forecast of μ in the linearized version of the model typically exceeds 0.99. In this sense, the solution is internally consistent.

Case	R^2	MFCE
Truth, approximated		
Stochastic π	0.9907 (−)	0.031% (−)
Stochastic λ	0.966 (−)	1.37% (−)
Stochastic π, λ	0.94 (−)	2.5% (−)
Linear, consistent		
Stochastic π	0.9908 (−)	0.3954% (−)
Stochastic λ	0.6958 (−)	0.7289% (−)
Stochastic π, λ	0.7032 (−)	1.707% (−)
Linear using DG truth		
Stochastic π	0.94 (−)	1.52% (−)
Stochastic λ	0.82 (−)	1.339% (−)
Stochastic π, λ	0.948 (−)	1.78% (−)

Table 4: Different approximation strategies

The first column shows the R^2 of a regression of output from non-adjustable plants on an intercept and the first moment, μ only. The second column reports the maximum forecast error from such a regression.

model with the RBC model in our environment. Given that higher order moments matter in the planner's optimization problem, it is natural to conjecture that the non-convexities also matter for aggregate moments.

Table 5 presents standard aggregate moments for a number of cases. These are the traditional macroeconomic moments: the correlations of output (y), consumption (c), investment (i) and TFP ($\tilde{A}$). Here the TFP measure is the one constructed from the data as if plants were homogeneous, i.e. mis-measured TFP. The serial correlations of consumption and output as well as relative standard deviations are reported, too.

The rows are the various cases explored before, using the baseline parameters. The last two rows "RBC" and "RBC QAC" are the standard stochastic growth model with productivity shocks, without and with quadratic adjustment costs. Here the productivity shocks come from fitting an AR(1) process to the mismeasured TFP series, $\tilde{A}$, generated by the stochastic (π, λ) case. We obtain an AR(1) parameter $\rho_{\tilde{A}} = 0.9183$ and standard deviation of the residual $\sigma_{\tilde{A}} = 0.0132$. The quadratic adjustment costs take the form $\frac{\gamma}{2}(\frac{I}{K})^2 \cdot K$. The estimate of $\gamma = 0.397$ was obtained by minimizing the distance between the moments of the model with stochastic (π, λ) and the RBC model with adjustment costs. The moments

Case	$C(y,c)$	$C(y,i)$	$C(y,\tilde{A})$	$C(i,\tilde{A})$	$C(i,c)$	ρ_c	ρ_i	$\frac{\sigma_c}{\sigma_i}$	$\frac{\sigma_c}{\sigma_y}$
Frictions									
stochastic A	0.857 (–)	0.785 (–)	0.978 (–)	0.838 (–)	0.35 (–)	0.903 (–)	0.801 (–)	0.359 (–)	0.860 (–)
stochastic π	0.826 (–)	0.691 (–)	0.963 (–)	0.727 (–)	0.164 (–)	0.855 (–)	0.779 (–)	0.381 (–)	0.951 (–)
stochastic λ	0.599 (–)	0.725 (–)	0.963 (–)	0.775 (–)	−0.118 (–)	0.685 (–)	0.573 (–)	0.253 (–)	0.898 (–)
stochastic π,λ	0.721 (–)	0.811 (–)	0.957 (–)	0.888 (–)	0.178 (–)	0.809 (–)	0.747 (–)	0.249 (–)	0.770 (–)
RBC	0.693 (–)	0.744 (–)	0.96 (–)	0.848 (–)	0.03 (–)	0.795 (–)	0.585 (–)	0.276 (–)	0.868 (–)
RBC QAC	0.77 (–)	0.769 (–)	0.961 (–)	0.846 (–)	0.185 (–)	0.833 (–)	0.60 (–)	0.303 (–)	0.847 (–)

Table 5: Endogenous Capital Accumulation - Macroeconomic Moments

Results from 1000 simulations are reported with standard deviations in parentheses below. Here $C(x,y)$ are correlations, ρ_x is an autocorrelation and σ_x is a standard deviation. The variables are: output (y), consumption (c), investment (i) and mis-measured TFP ($\tilde{A}$).

were: $C(y,c), C(y,i), \rho_c$ and ρ_i.

All of the models match the standard business cycle properties of positively correlated movements in consumption, investment and output. All of these variables move with (mis-measured) TFP. And these aggregate moments are all positively serially correlated.

Comparing the case of stochastic (π,λ) and the RBC model with adjustment costs, the moments are very close. Investment though is less serially correlated and less correlated with consumption in the RBC QAC treatment. This is interesting since a researcher would interpret the data as generated by a model with TFP shocks even though A is actually constant. That is, the researcher would misinterpret the actual $\tilde{A}$ shocks as variations in A.

6 Conclusion

The goal of this paper was to understand the productivity gains from capital reallocation in the presence of frictions. To study this we have looked at the optimization problem of a planner facing frictions in capital accumulation and shocks to productivity, adjustment costs and the distribution of plant specific shocks.

The heterogeneity in plant-level productivity provides the basis for reallocation. The frictions in adjustment prevent the full realization of these gains.

There are two key findings in this paper. The first is the cyclical behavior of reallocation. When shocks to either adjustment frictions or the distribution of plant-level shocks are present, then reallocation is procyclical. In fact, even if there are no direct shocks to TFP, the reallocation process creates fluctuations in output and investment. These effects are not present when the only shock is to TFP.

Second, in some, though not all environments, the plant-level covariance of capital and profitability shocks matters for characterizing the planner's solution. This is important for a few reasons. First, it is indicative of state dependent gains to reallocation. Second, our economy is an example of one where moments other than means are needed in the planner's problem.

Appendix

The appendix describes our method of solving the planner's problem. The approach taken for characterizing the law of motion for the joint distribution, Γ, is described in the text. Here we focus on the planner's choice of capital in the reallocation process.

Any vector of capital allocated across adjustable plants $k(\varepsilon)$ will have associated values for μ_a and ϕ_a. Create a grid for potential vectors $k(\varepsilon)$. To so so, define two benchmarks for the planners decision regarding the allocation of capital across those plants that are in F^A. Define k^{MAX} as the vector where marginal products are equalized across plants. This vector was found in (7) for the frictionless benchmark case above. In the presence of Calvo adjustment costs, the planner will not reallocate more capital between plants than under the allocation rule k^{MAX}, but possibly less. The second benchmark will be called k^{MIN} and is simply the case where capital is equally distributed across adjustable plants (i.e. no reallocation). The idea behind this procedure is that the planner will choose a vector $k(\varepsilon)$ which is *between* k^{MAX} and k^{MIN}, meaning that the planner will reallocate some capital between plants, but not as much as under the frictionless benchmark. We consider convex combinations of k^{MAX} and k^{MIN}.

Define a variable m, that takes values between zero and one and determines a potential vector of $k(\varepsilon)$'s as follows: $k_m = m \cdot k^{\text{MAX}} + (1-m) \cdot k^{\text{MIN}}$. For each k_m compute $\mu_m = E(k_m(\varepsilon)^\alpha)$ and $\phi_m = Cov(\varepsilon, k_m(\varepsilon)^\alpha)$ characterizing this vector. This allows the calculation of output associated with m. The planner optimizes over m and this translates into μ_m, ϕ_m.

To check the robustness of this procedure start from a model with the baseline parameters without any exogenous shocks. It turns out that the planner chooses $m = 0.9508$, which means that the optimal vector $k(\varepsilon) = 0.9508 \cdot k^{\text{MAX}} + 0.0492 \cdot k^{\text{MIN}}$, so capital reallocation is about 5% lower compared to the frictionless benchmark. In order to see how good of an approximation the decision rule 'm' is, we apply the following procedure.

We work directly with the planner's value of the steady state (SS) allocation. The simplified version of the value function has only two states, μ_n and ϕ_n, so there will be a value $V(\mu_n^{SS}, \phi_n^{SS})$ associated to the steady state. This value is equal to forever receiving the output associated with the amount of reallocation 'm' times the fraction of adjustable plants, plus the output associated with the SS state vector times the fraction of non-adjustable plants.

$$V(\mu^{SS}, \phi^{SS}) = \frac{\int_{\varepsilon \in F^A} \varepsilon k(\varepsilon)^\alpha f(\varepsilon) d\varepsilon \ \ + (1-\pi)(E(\varepsilon)\mu^{SS} + \phi^{SS})}{1-\beta} \tag{28}$$

The planner can now choose any allocation of capital across plants. This allocation implies a mapping into the values of μ_n and ϕ_n. The planner will be allowed to choose the allocation that maximizes the expression for $V(\mu^{SS}, \phi^{SS})$ above. Being bound to the same grid, the resulting vector is identical to the one previously found. We now perturb this vector in order to find profitable deviations that keep the aggregate capital stock constant. The perturbation adds a random vector with mean zero to the k-vector that maximized (27) given the grid. If the resulting vector produces a higher lifetime utility, the k-vector is updated accordingly. This procedure is repeated 1,000,000 times. The results show that our grid for m comes extremely close to the optimal solution. Although profitable deviations *are* possible, they remain very small: the difference in output is around 0.01%.

References

Bartelsman, E., J. Haltiwanger, and S. Scarpetta (2006): "Cross Country Differences in Productivity: The Role of Allocative Efficiency," Available at http://www.econ.umd.edu/ ˜haltiwan/Creative_destruction_BHS.pdf.

Basu, S., and J. Fernald (1997): "Returns to Scale in U.S. Production: Estimates and Implications," *Journal of Political Economy*, 105, 249–83.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77(3), 623–685.

Calvo, G. (1983): "Staggered prices in a utility-maximizing framework," *Journal of monetary Economics*, 12(3), 383–398.

Den Haan, W. J. (2010): "Assessing the accuracy of the aggregate law of motion in models with heterogeneous agents," *Journal of Economic Dynamics and Control*, 34(1), 79–99.

Eisfeldt, A. L., and A. A. Rampini (2006): "Capital reallocation and liquidity," *Journal of Monetary Economics*, 53(3), 369–399.

Gilchrist, S., J. Sim, and E. Zakrajsek (2010): "Uncertainty, Financial Frictions and Investment Dynamics'," Available at.

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

Gourio, F., and J. Miao (2006): "Firm Heterogeneity and the Long-Run Effects of Dividend Tax Reform," Available at http://people.bu.edu/fgourio/divGEpaper.pdf.

Hsieh, C.-T., and P. Klenow (2007): "Misallocation and Manufacturing TFP in China and India," Available at http://www.klenow.com/MMTFP.pdf.

Kehrig, M. (2011): "The Cyclicality of Productivity Dispersion," The University of Texas at Austin, Job Market Paper.

Khan, A., and J. Thomas (2003): "Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities Matter?," *Journal of Monetary Economics*, 50, 331–60.

——— (2006): "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics," Available at http://www.juliathomas.net/.

Krusell, P., and A. Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106(5), 869–96.

Lucas, R. E. (1978): "On the Size Distribution of Business Firms," *Bell Journal of Economics*, 9(2), 508–523.

Olley, S. G., and A. Pakes (1996): "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64(6), 1263–1297.

Restuccia, D., and R. Rogerson (2007): "Policy Distortions and Aggregate Productivity with Heterogeneous Plants," Available at http://www.nber.org/papers/w13018.

Sveen, T., and L. Weinke (2005): "New perspectives on capital, sticky prices, and the Taylor principle," *Journal of Economic Theory*, 123(1), 21–39.

Thomas, J. (2002): "Is Lumpy Investment Relevant for the Business Cycle?," *Journal of Political Economy*, 110, 508–34.