

Capital reallocation and the cyclicality of aggregate productivity

RUSSELL W. COOPER

Department of Economics, European University Institute and NBER

IMMO SCHOTT

Department of Economics, Université de Montréal

Capital reallocation between firms is procyclical and leads to variations in measured aggregate productivity. In this paper, we ask how much of the cyclical variation in measured productivity is the consequence of capital reallocation. We build a heterogeneous-firm model to study the effects of exogenous shocks to total factor productivity (TFP) and to the costs of reallocation. These shocks cause an endogenous cyclical variation of measured aggregate productivity. Only a model driven by exogenous TFP shocks is able to generate both data-consistent cyclical movements in reallocation and sizeable variations in measured aggregate productivity. We find that capital reallocation does not play a major role in amplifying aggregate productivity variations over the business cycle.

KEYWORDS. Reallocation, productivity.

JEL CLASSIFICATION. E22, E23, E24.

1. INTRODUCTION

In a given year, about 6% of the existing capital stock of U.S. public companies is reallocated between firms through sales of existing capital and acquisitions. This capital reallocation is sizeable, adding up to almost a quarter of total investment. In an average year between 1972 and 2018, 43% of firms sell part of their existing capital, while 29% make an acquisition (12% do both).

The existing literature establishes a number of key facts about capital misallocation, which are informative about the gains from reallocation. First, from Restuccia and Rogerson (2008), Hsieh and Klenow (2009), and others, misallocation can explain long-run differences in economic performance. Second, there are productivity gains from reallocation, as established in Olley and Pakes (1996), Maksimovic and Phillips (2001), and Foster, Haltiwanger, and Syverson (2008) among others. Finally, Eisfeldt and Rampini (2006) documented that the capital reallocation process is procyclical, that is, more capital is reallocated during times when aggregate output is high.

Russell W. Cooper: russellcoop@gmail.com

Immo Schott: immoschott@gmail.com

We are grateful to the referees for many helpful suggestions. We thank Nick Bloom, Jesus Bueren, Thorsten Drautzburg, Michael Elsby, John Haltiwanger, Matthias Kehrig, Andrea Lanteri, Ryan Michaels, Lukas Nord, Sophie Osotimehin, Felipe Saffie, and Thomas Winberry, as well as to seminar participants for comments and suggestions. Cooper thanks the NSF under grant 0819682 for financial support. A replication file is posted (Cooper and Schott (2023)).

The correlation highlighted by Eisfeldt and Rampini (2006) is suggestive: capital reallocation may contribute to fluctuations in measured aggregate productivity. If so, then observed cyclical variations in the Solow residual might be explained as an outcome of the reallocation process. More capital reallocation leads to a better allocation of resources across firms, which implies higher measured aggregate productivity. This paper explores this conjecture. We ask: How much of the observed cyclical variation in aggregate productivity can be explained by capital reallocation?

The framework for analysis is a heterogeneous firm model in which the movement of capital across producers entails adjustment frictions. This implies productivity gains from reallocation because marginal products are not equalized. Changes in aggregate conditions lead to endogenous variations in capital reallocation. This causes fluctuations in the Solow residual, or measured aggregate productivity.

We solve the model as a planner's problem in an economic environment that isolates the role of frictions to the capital reallocation process. We study a tractable, yet rich model of capital adjustment frictions, which can be interpreted as costs of information acquisition. These costs generate state independent firm-level adjustment rates that match the distribution of reallocation rates and the amount of inaction observed in the data. Although this setup entails higher order moments of the joint cross-sectional distribution of productivity and capital, our approach allows for an analytical characterization of the solution.¹

The model is calibrated to match both the key microeconomic and macroeconomic facts about the capital reallocation process. Specifically, the steady state of the model exactly matches the empirical targets on the fraction of reallocating firms, the average amount of capital reallocation, and the dispersion in revenue-based total factor productivity. As indicated by the distribution of reallocation rates in Figure 1, a key feature of firm-level data is a large amount of heterogeneity in reallocation rates.² The model captures this heterogeneity and quantifies its importance for cyclical changes in aggregate productivity.

We then use this model as a foundation for the study of cyclical reallocation and productivity. Two principal sources of fluctuations are introduced: (i) shocks to aggregate TFP and (ii) shocks to the distribution of adjustment costs. These shocks are propagated by the evolution of the cross-sectional distribution of productivity and capital, a state variable of the model. The parameters for these stochastic processes are estimated using empirical moments on the cyclicity of capital reallocation, as well as the time-series variation in the cross-sectional dispersion of reallocation rates. This choice of moments ensures that the key macroeconomic characteristic of aggregate reallocation is generated by the model economy, while at the same time reproducing realistic cyclicity in the *amount* of capital reallocation.

¹In the baseline model, aggregate capital is fixed and labor can be adjusted without frictions. We discuss the model assumptions in more detail in Section 3.2 below. An extension in Section 4.4.2 adds capital accumulation.

²The figure includes only positive capital reallocation rates. The fraction of firms reallocating is about 60% each year.




FIGURE 1. Histogram of firm-level reallocation rates.

We have two main findings: First, the model economies driven by exogenous variations in either aggregate productivity or the costs of reallocation can induce movements in capital reallocation that match the micro and macro data on reallocation and generate realistic business cycles. In both economies, the endogenous movement of capital across heterogeneous firms matches the cyclicity of reallocation and creates an endogenous, procyclical Solow residual. However, the economy driven by variations in capital adjustment costs only generates about 4.5% of the observed variation in aggregate productivity. We conclude that the correlation between productivity and reallocation stems from the effects of exogenous productivity shocks on reallocation. Our second finding is that in the model driven by TFP shocks, only about 3% of the standard deviation of the Solow residual comes from the resulting capital reallocation. This is not to say that reallocation is unimportant for longer-run economic phenomena. But at the business cycle frequency, reallocation of capital is not a driving force.

Related literature

A large body of work has studied the process of capital reallocation empirically (e.g., Olley and Pakes (1996), Maksimovic and Phillips (2001), Ramey and Shapiro (2001)). A main conclusion is that the reallocation process is productivity-enhancing (Foster, Haltiwanger, and Syverson (2008), Osotimehin (2019)). Such findings have motivated the question to what extent frictions to the reallocation process can explain the observed productivity differences both across and within countries. The common idea of this literature is that frictions in the reallocation process lead to the misallocation of factors of production (relative to a frictionless benchmark) and affect aggregate productivity. Removing these frictions, which can either take the form of unspecified “wedges” (Restuccia and Rogerson (2008), Hsieh and Klenow (2009)) or specific policies (Guner, Ventura, and Xu (2008), Kaymak and Schott (2019)) is found to entail large productivity

gains, especially when the frictions are positively correlated with firm-level productivity (Bartelsman, Haltiwanger, and Scarpetta (2013)). Importantly, these papers study long-run productivity gains. In this paper, we investigate the conjecture that at least some of the fluctuations in aggregate productivity at the business cycle frequency could be the consequence of the dynamics of capital reallocation.

Our analysis builds upon Eisfeldt and Rampini (2006), who established two important empirical regularities: Among publicly listed U.S. firms, capital reallocation is strongly procyclical, while measures of productivity dispersion, interpreted as “benefits” to reallocation, are countercyclical or acyclical.³ The results in Eisfeldt and Rampini (2006) suggest an important role for cyclical movements in the costs of reallocation. Differently from Eisfeldt and Rampini (2006), our economic environment includes a joint distribution of capital and idiosyncratic productivity. This allows us to perform a quantitative analysis using a model economy that replicates both macroeconomic regularities and the underlying microeconomic heterogeneity. Further, our analysis highlights the role of nonconvex adjustment costs, allowing us to match moments on both the extensive and intensive margins of firm-level changes in capital.

Lanteri (2018) studies capital reallocation in a model with an explicit market for secondary capital to generate endogenous partial irreversibility. He presents evidence for a procyclical secondary market price of capital and generates procyclical capital reallocation in a model with aggregate productivity shocks. Lanteri (2018) focuses on the business cycle properties of a DSGE model with an endogenous resale price of capital. We focus solely on capital reallocation with an emphasis on cross sectional properties of the economy. Our aim is to construct a model that matches both the cyclical character of aggregate capital reallocation, as well as the heterogeneity in microlevel reallocation behavior. This ensures that both the overall amount of capital reallocation, and the dispersion in reallocation rates are data-consistent. This is important given our focus on understanding sources of exogenous variation that explain sizeable variations in aggregate productivity and match the reallocation moments. One finding in Lanteri (2018) is that an economy subject only to aggregate productivity shock generates little endogenous feedback from capital reallocation. Our results confirm this result. We go further by investigating other sources of exogenous variation, particularly variations in the cost of reallocation as stressed in Eisfeldt and Rampini (2006) and show that these alternative causes of business cycles either fail to match the reallocation moments or fail to generate sizeable movements in measured aggregate productivity.

Chen and Song (2013) study capital reallocation and the resulting productivity implications of news shocks. In their model, news shocks loosen financial constraints, which leads to capital reallocation from unconstrained to constrained firms, causing variations in measured total factor productivity. Our analysis does not include capital market imperfections and focuses on others sources of aggregate fluctuations.

Finally, in comparison to most quantitative dynamic general equilibrium models, we do not rely on a computationally heavy quantitative solution in the tradition of Krusell

³Countercyclical productivity dispersion is central to many studies and this empirical regularity has often been replicated, for example, in Bloom (2009), Kehrig (2015), and Bachmann and Bayer (2014).

and Smith (1998). Our setup allows for an analytical characterization of the solution, which entails higher-order moments of the cross-sectional distribution of productivity and capital. We are able to analyze to what extent quantitatively plausible variations in aggregate conditions are able to generate capital reallocation and affect measured aggregate productivity.

2. REALLOCATION AND PRODUCTIVITY

To fix basic ideas, consider an economy without aggregate shocks.⁴ The economy is composed of a representative household and a mass of firms distinguished by productivity and adjustment costs. We first discuss the basic environment. An important derivation links the Solow residual to the cross sectional distribution of capital. Then the stationary equilibrium is computed and matched to a subset of moments.

2.1 Environment

The production function at any firm is

$$y(k, \varepsilon) = A\varepsilon k^\alpha, \quad (1)$$

where A is (fixed) aggregate TFP, and k is the capital used at the firm with idiosyncratic productivity ε .⁵ The aggregate capital stock is fixed at K , allowing us to focus on reallocation alone.⁶

Idiosyncratic productivity shocks are persistent. Every period, with probability $1 - \rho_\varepsilon$ a production site draws a new level of ε from a log-normal distribution. This new draw is unconditional on ε_{t-1} . With the counterprobability ρ_ε the productivity level remains unchanged:

$$\varepsilon_t = \begin{cases} \varepsilon_{t-1}, & \text{with probability } \rho_\varepsilon \\ \ln \mathcal{N}(0, \sigma_\varepsilon), & \text{with probability } 1 - \rho_\varepsilon \end{cases}. \quad (2)$$

The invariant distribution of ε is denoted $f(\varepsilon)$.

We assume $\alpha < 1$ as in Lucas (1978). The assumption of diminishing returns to scale implies that the allocation of capital across production sites is nontrivial. There are gains to allocating capital to high productivity sites but as $\alpha < 1$, capital is also spread across firms.

Following the realizations of firm-specific ε , capital can be reallocated across production sites. Capital reallocation is costly. Our approach is to specify and estimate a minimalist adjustment cost structure. Specifically, in order to learn about the firm-specific state (k, ε) before reallocation and production take place, a fixed adjustment cost F must be paid. The adjustment cost F is independently and identically distributed

⁴Aggregate shocks are added in Section 4 for the study of cyclical reallocation.

⁵The model does not distinguish between plants and firms. Labor and other inputs are not made explicit. One interpretation is that these inputs have no adjustment costs and are optimally chosen each period, given the state.

⁶Section 4.4.2 adds capital accumulation to the model.

across time and firms with a CDF given by $G(F)$. As in Eisfeldt and Rampini (2006), these costs can be interpreted as information frictions rather than physical adjustment frictions.

Our emphasis on information frictions, which generates a state independent adjustment rate, has three distinct advantages.⁷ First, combining costs of adjustment and information acquisition makes the analysis considerably more tractable. In fact, many of our results are analytic. Second, the specification cleanly separates the costs and benefits of reallocation. Third, many models of costly capital adjustment have difficulty matching small investment rates.⁸

Our setup has the potential to generate both inaction region as well as a distribution of the intensive margin of capital adjustments. Small, positive adjustments can occur because after paying the information cost, capital adjustment might be minimal if the current level of capital is sufficiently close to the optimal level. Section 3.2 discusses the role of these assumptions for the generality of our results in greater detail.

2.2 Optimal allocations

In the presence of reallocation costs, the choice problem of a social planner is

$$V(\Gamma(k, \varepsilon)) = \max_{\pi, \tilde{k}_a(k, \varepsilon)} u(c) + \beta E_{\Gamma|\Gamma} V(\Gamma'(k, \varepsilon)) \quad (3)$$

subject to

$$y = \int_{(k, \varepsilon) \in a} AK^\alpha \varepsilon \tilde{k}_a(k, \varepsilon)^\alpha d\Gamma(k, \varepsilon) + \int_{(k, \varepsilon) \in n} AK^\alpha \varepsilon \tilde{k}_n(k, \varepsilon)^\alpha d\Gamma(k, \varepsilon), \quad (4)$$

$$K = \int_{(k, \varepsilon) \in a} \tilde{k}_a(k, \varepsilon) d\Gamma(k, \varepsilon) + \int_{(k, \varepsilon) \in n} \tilde{k}_n(k, \varepsilon) d\Gamma(k, \varepsilon), \quad (5)$$

$$c = y - K \int_0^{F(\pi)} F dG(F). \quad (6)$$

In this optimization problem, the state variables of the planner are the joint distribution over firm-level productivity and capital prior to reallocation, denoted by $\Gamma(k, \varepsilon)$, explained below. There are two controls in (3). The planner determines π , the fraction of firms that will be adjusted, and chooses how to allocate capital among adjusting firms. The adjustment status of a firm is given by $j \in \{a, n\}$, where a stands for “adjusting,” while n stands for “not adjusting.” This status is determined by whether or not a firm’s fixed cost of adjusting is paid. Let $\tilde{k}_j(k, \varepsilon)$ for $j \in \{a, n\}$ denote the fraction of aggregate capital allocated to a firm with adjustment status j that enters the period with capital k and productivity ε . The capital of firms in the set of adjusters is optimally set by the planner

⁷This parallels the debate in the pricing literature between information and menu cost frictions as the source of sticky prices. Klenow and Willis (2007) review these models and provide some evidence in favor of the information cost models.

⁸For example, Cooper and Haltiwanger (2006) match the extremes of the investment rate distribution but generate no small adjustments.

to a level $\tilde{k}_a(k, \varepsilon)$. The capital of firms in the group of nonadjusters remains unchanged, that is, $\tilde{k}_n(k, \varepsilon) = k$.

Here, output is simply the sum of the production from adjusting and nonadjusting firms, as in (4), while aggregate capital is the sum of capital in the two types of firms, as in (5). In the resource constraint (6), the adjustment cost are linked to the fraction of adjusting firms π through the CDF of adjustment costs.⁹ Specifically, given π , the planner selects firms starting with the lowest adjustment costs until the desired fraction π of firms are adjusted. Through this process, the maximal cost incurred is denoted $F(\pi)$ and given implicitly by $\pi = G(F)$. Once the maximal adjustment cost is determined, the total amount paid is the integral over the distribution of adjustment costs up to $F(\pi)$, as in the last term of (6).

As a benchmark, consider the planner's problem in the absence of adjustment costs. The objective is to maximize output (consumption) through an optimal allocation of capital. The cross-sectional distribution is no longer a state variable. The resulting first-order condition is

$$\tilde{k}_a(k, \varepsilon) = K \frac{\varepsilon^{\frac{1}{1-\alpha}}}{\int_{\varepsilon} \varepsilon^{\frac{1}{1-\alpha}} f(\varepsilon) d\varepsilon} \quad (7)$$

for all k, ε .¹⁰ The optimal capital allocation equalizes marginal products of capital across firms. For adjusters, this allocation depends on ε but, as is clear from the right side of (7), is independent of k .

Adjustment costs imply that the reallocation problem is dynamic. They create two distinct dimensions linked to the capital reallocation problem, an intensive and an extensive margin. The extensive margin relates to the fraction of adjusting firms π , while the intensive margin is controlled by the planner through the assignment function $\tilde{k}_a(k, \varepsilon)$. Both instruments influence the distribution of capital across firms but imply different costs. Choosing a higher π enlarges the set of adjusting firms. This has the benefit of increasing output today but increases the total adjustment costs. Choosing an allocation of capital $\tilde{k}_a(k, \varepsilon)$ that reduces the dispersion in marginal products of capital also increases output today. Here, costs arise due to future adjustment costs. Given that a fraction $1 - \rho_\varepsilon$ of firms will receive a new productivity draw next period, a fraction of firms will not be at their optimal k in the future. This entails larger future extensive margin adjustments and higher future adjustment costs. For this reason, $\tilde{k}_a(k, \varepsilon)$ will differ from the frictionless capital allocation satisfying (7).

We define reallocation of capital as the fraction of total capital that is moved between adjusting firms within a period. As $\tilde{k}_a(k, \varepsilon)$ denotes the post-reallocation capital stock of adjusting firms with initial capital k and productivity ε , firm-level reallocation level would be $r(k, \varepsilon) = |\tilde{k}_a(k, \varepsilon) - k|$. Aggregating over all adjusting firms, the aggregate reallocation level is

$$R \equiv 0.5 \int_{(k, \varepsilon) \in a} r(k, \varepsilon) d\Gamma(k, \varepsilon). \quad (8)$$

⁹As in Cooper and Haltiwanger (2006), the adjustment cost is proportional to K .

¹⁰Derivations can be found in Appendix B.

The multiplication by 0.5 avoids double counting flows between adjusting firms.

2.3 Aggregate productivity

For this economy, there exists a fundamental link between productivity and the assignment of capital to firms. This link lies at the heart of our analysis of reallocation and aggregate productivity. Aggregate output in (4) can be rewritten as

$$y = AK^\alpha [\pi(\mu_a + \phi_a) + (1 - \pi)(\mu_n + \phi_n)] = \tilde{A}K^\alpha, \quad (9)$$

where, for $j = \{a, n\}$, define $\mu_j \equiv \mathbb{E}(\varepsilon \tilde{k}_j(k, \varepsilon)^\alpha)$ as the product of average productivity and average effective capital, while $\phi_j \equiv \text{Cov}(\varepsilon, \tilde{k}_j(k, \varepsilon)^\alpha)$ is the covariance between the two terms.¹¹

The values μ_n and ϕ_n pin down output among nonadjusting firms and are determined by the current state, the joint distribution of productivity and capital Γ . Output among adjusting firms is controlled through the assignment function $\tilde{k}_a(k, \varepsilon)$. This assignment of capital to adjusting firms implies values for the mean and covariance terms μ_a and ϕ_a . Finally, the choice of a fraction of adjusting firms π determines the relative mass of adjusters and nonadjusters.

Researchers interested in measuring TFP from the aggregate data will typically uncover $\tilde{A}$, the Solow residual, rather than A . From (9), there are two factors which influence $\tilde{A}$. The first one is A . The influence of A , aggregate TFP, is direct and has been central to many studies of aggregate fluctuations. The second effect on $\tilde{A}$ comes from endogenous changes in the allocation of capital to production sites. Both changes in the fraction of adjusting firms π (the extensive margin), as well as changes in the target allocation of capital $\tilde{k}_a$ (the intensive margin) impact the Solow residual. The aggregate shocks we consider below will lead to endogenous changes in capital reallocation, thereby impacting the Solow residual through reallocation.

2.4 Joint distribution of capital and productivity

In the presence of reallocation frictions, the state of the economy includes the joint distribution of firm specific capital, k , and the idiosyncratic shock, ε , denoted as $\Gamma(k, \varepsilon)$. This distribution determines the level of output that is produced at nonadjusting firms. Consequently, when making reallocation decisions the planner needs to forecast Γ' . The evolution of this distribution reflects the persistence in the idiosyncratic shocks as well as the choices about the fraction of adjusting firms π and the assignment of capital among adjusting firms $\tilde{k}_a(k, \varepsilon)$.

An advantage of our setup is that although the model is rich due to firm heterogeneity and nonconvex adjustment costs, it allows for an exact analytical representation

¹¹This builds on Olley and Pakes (1996). To understand the derivation, suppose all firms reallocate, $\pi = 1$. From $y(k, \varepsilon) = A\varepsilon k^\alpha$, $y = \int_{(k, \varepsilon)} AK^\alpha \varepsilon \tilde{k}_a(k, \varepsilon)^\alpha d\Gamma(k, \varepsilon) = AK^\alpha \int_{(k, \varepsilon)} \varepsilon \tilde{k}(k, \varepsilon)^\alpha d\Gamma(k, \varepsilon) = AK^\alpha (\mathbb{E}(\varepsilon \tilde{k}(k, \varepsilon)^\alpha) + \text{Cov}(\varepsilon, \tilde{k}(k, \varepsilon)^\alpha))$. Equation (9) extends this logic to include adjusting and non-adjusting firms.

of the joint distribution and its evolution. Different from the literature following Krusell and Smith (1998) we are not bound by computational limitations to a numerical approximation of this joint distribution. Instead of keeping track of the infinite-dimensional joint distribution Γ , we retain μ_n and ϕ_n in the state vector of (3). Together with the choices π and $\tilde{k}_a(k, \varepsilon)$, this pins down contemporaneous output in (9). Importantly, the higher-order moment, ϕ_n is part of the state of economy. This moment represents the dependence of output on the assignment of capital to firms. Dynamics in the cross-sectional joint distribution of capital and productivity will induce time-series variation in the Solow residual even in the absence of serially correlated aggregate shocks. The moments μ_n and ϕ_n are essential for capturing these dynamics in the analysis that follows.

The law of motion for the distribution Γ can be written as a convex combination of the moments from the adjusting and nonadjusting firms, weighted respectively by π and $1 - \pi$:¹²

$$\mu'_n = (1 - \pi)\mu_n + \pi\mu_a, \quad (10)$$

$$\phi'_n = \rho_\varepsilon \cdot ((1 - \pi)\phi_n + \pi\phi_a). \quad (11)$$

Note that in the absence of aggregate shocks, this economy has a stationary joint distribution of capital and productivity, given by

$$\mu_n = \mu_a, \quad (12)$$

$$\phi_n = \phi_a \cdot \frac{\pi}{\pi + \frac{1 - \rho_\varepsilon}{\rho_\varepsilon}}. \quad (13)$$

From (9), this implies that steady-state consumption is given by

$$c = AK^\alpha \left(\mu_n + \frac{\phi_n}{\rho_\varepsilon} \right) - K \int_0^{F(\pi)} F dG(F). \quad (14)$$

These characterizations reflect the tradeoffs the planner is facing. Spreading capital more evenly across firms increases μ_n but lowers the covariance with productivity ϕ_n . Increasing the fraction of adjusting firms π increases the steady-state covariance between sites' productivity and capital from (13), but raises the amount of adjustment costs paid.

3. QUANTITATIVE ANALYSIS: STEADY STATE

In this section, we estimate a quantitative version of the model developed in Section 2. The goal is to match key characteristics of the microeconomic adjustment behavior of firm-level changes in capital. For this, the steady-state properties of the model are matched with those of the data to set a couple of key parameters characterizing firm-level choices. Details of our solution method can be found in Appendix C.

¹²See Appendix B for derivations.

TABLE 1. Parameters.

Parameter	Meaning	Value	Source
β	Discount factor	0.962	Annual $r = 4\%$
γ	Risk aversion	1	Log utility
α	Curvature of revenue function	0.810	Compustat
ρ_ε	Persistence of ε	0.930	cf. Table 2
σ_ε	Std. dev. of ε	0.432	
$\bar{F}$	Adjustment cost upper bound	0.169	cf. Table 3
ρ_A	Persistence of A	0.655	
σ_A	Standard deviation of A	0.015	
$\rho_{\bar{F}}$	Persistence of $\bar{F}$	0.948	
$\sigma_{\bar{F}}$	Standard deviation of $\bar{F}$	0.044	

Note: Model parameters. The parameters in upper part of the table were preassigned. The parameters in the second part of the table were estimated for the stationary economy using SMM. The lower part of the table shows the parameters of the various aggregate shock processes for the model with aggregate fluctuations.

3.1 Parametrization

We begin by calibrating the steady state of the model described in Section 2.2. The parameters are chosen by targeting moments that characterize capital reallocation and the importance of idiosyncratic shocks.¹³

There are six parameters in the steady state of the model, summarized in the upper part of Table 1. The first two parameters describe the household’s discount factor and risk aversion and are set outside the model. To determine the returns to scale parameter α , we estimated the curvature of the revenue function using firm-level data and found a value of 0.81.¹⁴

The second part of Table 1 presents the results of a simulated method of moments (SMM) estimation. We estimate three parameters. The persistence and standard deviation of idiosyncratic shocks are denoted as ρ_ε and σ_ε . To parameterize the adjustment costs, we assume that $G(F)$ is uniform between zero and an upper limit denoted $\bar{F}$, as in Thomas (2002). The parameter vector $\Theta = (\rho_\varepsilon, \sigma_\varepsilon, \bar{F})$ is chosen to minimize the distance between data and the targeted model moments, as shown in the first part of Table 2.

While there is no one-to one mapping between parameters and moments, the data targets were chosen to be informative about the underlying parameters. The first two moments are the fraction of firms that adjust capital and the fraction of total capital that is reallocated in a given period. These moments are informative about the importance of the capital adjustment cost and about the persistence of idiosyncratic productivity shocks. If reallocation is less costly, capital in a larger fraction of firms will be reallocated, but a lower amount of reallocation is required each period. A higher persistence of idiosyncratic productivity implies lower reallocation rates and a lower fraction of adjusting firms.¹⁵ The third moment is the value-added weighted standard deviation of firm-level

¹³The data moments were computed using Compustat. Details can be found in Appendix A.
¹⁴We provide a robustness analysis with respect to this and other parameters in Appendix D.
¹⁵Note that, as stressed in Lanteri (2018), Compustat provides a measure of the dollar value of capital reallocation, so that changes in reallocation might be driven by prices rather than quantities. Our measures

TABLE 2. Moments from the stationary economy.

	Targeted			Untargeted			
	$R > 0$	R/K	σ_{TFPR}	ρ_{sales}	$\rho_{R/K}$	$c(\text{ARPK}, R/K)$	small R/K
Data	0.619	0.057	0.363	0.908	0.004	0.053	0.831
Model	0.620	0.057	0.363	0.940	-0.002	0.001	0.868

Note: The first three moments are targeted in the calibration. Those include (i) the time-series average of the fraction of firms with positive capital reallocation, (ii) the time-series average of the amount of capital reallocation (as a fraction of total capital), (iii) the time-series average of the value-added weighted standard deviation of firm-level revenue-based productivity (TFPR). The untargeted moments are (i) the serial correlation of the logarithm of firm-level sales, (ii) the serial correlation of firm-level reallocation rates, (iii) the correlation between the firm-level average revenue product of capital (ARPK) and reallocation rates, and (iv) the fraction of small adjustment rates. Details can be found in Appendix A.

productivity. We measure revenue-based productivity in the data using a semiparametric procedure following Olley and Pakes (1996). In the model, the standard deviation in firm-level productivity is governed by σ_ε . Additionally, a higher σ_ε implies that productivity is more dispersed among firms, leading to more reallocation.

Table 2 shows an exact fit between the targeted data and model moments. The model reproduces key features of the data on capital reallocation and productivity dispersion across firms.¹⁶ In particular, the model generates the large fraction of firms with positive capital adjustments, while simultaneously matching both the fraction of total capital that is being reallocated and the consequent firm-level dispersion in revenue-based productivity (TFPR).¹⁷

The model also fits a series of untargeted moments, shown in the last four columns of Table 2. The persistent productivity shocks imply a high serial correlation of firm-level (log) sales of over 0.90, while the state-independent nature of the adjustment costs implies almost zero serial correlation in plant-level capital adjustment rates. These remaining moments are discussed in the next section.

3.2 Firm-level capital adjustment

Despite its simplicity, the model is able to generate a microeconomic adjustment behavior that is in line with the empirical counterpart. First, there is a large amount of inaction. Almost 40% of firms do not reallocate capital in a given year. Second, as was shown in Figure 1, the empirical distribution of capital reallocation is characterized by many small adjustments. From the figure, most active firms only make small adjustments to their capital stock. We quantify this moment by measuring the fraction of firms that represent 10% of total capital reallocation in a given year. In the data, the time-series average of this number is 83.1%, pointing to a significant skewness in reallocation rates. This moment is shown in the last column of Table 2. The quantitative model matches

of reallocation are not subject to this concern because $R > 0$ is an extensive margin measure and R/K measures the fraction of reallocated capital.

¹⁶Aggregate capital reallocation costs amount to 2.3% of output in the calibrated model.

¹⁷In the data, it is TFPR that is measured. In the model, there is no difference between TFPR and aggregate productivity TFP.




FIGURE 2. Average adjustment rates by average revenue product of capital (ARPK).

this number very well, with the model counterpart being that 86.8% of capital adjustments are small in this sense. This success is a result of our formulation of the adjustment costs: after paying the information acquisition cost, capital adjustment might be minimal if the current level of capital is sufficiently close to the optimal level of capital given ε , leading to many small, positive adjustments.

In the model, firms’ adjustment behavior is independent of the individual firm state. We test this assumption by showing how adjustment rates in the data correlate to firms’ average revenue product of capital (ARPK), a measure of firm productivity. This is shown in Figure 2. Adjustment rates in the data are essentially independent of the average revenue product of capital. To the extent that differences in ARPKs are related to the gains to reallocation, say through the gap approach popularized in Caballero and Engel (1999), adjustment rates would be highest in the tails of the distribution in a state dependent adjustment model. This is not the case in our sample.

We summarize this through the correlation between firm-level ARPK and capital reallocation rates. From Table 2, this correlation is very low in the data, a feature shared by the model.

3.3 Reallocation and aggregate productivity

Before we study the role of capital reallocation over the business cycle, we make use of the tractability of our model setup to show that changes in the economy’s allocative efficiency have potentially large effects for aggregate productivity in the steady state. Capital reallocation across heterogeneous firms generates productivity gains due to frictions. We quantify those productivity gains from reallocation on the Solow residual, which is defined as

$$\tilde{A} = A \cdot (\pi \cdot (\mu_a + \phi_a) + (1 - \pi) \cdot (\mu_n + \phi_n)). \tag{15}$$

There are two reasons why reallocation can be productivity enhancing. First, conditional on adjustment, there is the intensive margin. After having paid the adjustment costs, the allocation of capital among adjusting firms $\tilde{k}_a(k, \varepsilon)$ is chosen to optimally solve the tradeoff between exploiting differences in productivity and decreasing returns. The resulting allocation is characterized by a higher covariance between productivity and capital ϕ_a than for nonadjusting firms. In our calibrated economy, ϕ_a is 12% higher than ϕ_n .¹⁸ Among nonadjusting firms, changes in idiosyncratic productivity may have caused the capital stock to be out of sync with its optimal level. As a result, the dispersion in average products of capital is an order of magnitude larger among nonadjusting firms. Second, there is the extensive margin. From (15), a higher fraction of adjusting firms π reallocates aggregate activity toward adjusting firms, thereby increasing aggregate output at the margin.

This analysis suggests that in an economy's steady state there are potentially large gains from an improved factor allocation. Figure 3 shows the marginal effect of changes in the fraction of active firms, π , on the steady-state level of the Solow residual, $\tilde{A}$. To do so, we exogenously vary the level of π and solve for the optimal assignment of capital to adjusting firms conditional on the new fraction of adjusters. From (15), the resulting level of $\tilde{A}$ is different, not only because of changes in the fraction of adjusting firms π , but also because the reoptimized capital assignment to adjusting firms affects μ_a and ϕ_a , and—because of the persistence of productivity and the law of motion of the steady-state covariance term ϕ —also the output of nonadjusting firms.

Figure 3 makes clear that for low levels of π , the marginal effect on measured aggregate productivity is very large because of substantial differences in marginal productivities among adjusting and nonadjusting firms. For example, our results imply a 12.7% decrease in aggregate productivity if, say, due to higher adjustment frictions, π were to be reduced to 25% of the level currently observed among U.S. firms. There are possibly large losses from reductions in capital reallocation.

The steady-state value of π that resulted from the calibration of the model is indicated by the red star in Figure 3. Our calculations imply that there are few potential long-run gains from further facilitating capital reallocation in the U.S. economy. The key to understanding this result is that there are two dimensions to capital reallocation, an average and a marginal effect. If a large number of firms selects into capital adjustment, this significantly increases capital reallocation and measured aggregate productivity. Yet at the same time, the marginal effect of further increases in reallocation on productivity is very small. The marginal effect at the red dot is small because most of the long-run gains to reallocation have already been exhausted.

This finding is important for understanding the effects of reallocation on productivity over the business cycle. That analysis requires the next step of introducing aggregate shocks into our framework.

¹⁸In the steady state $\mu_a = \mu_n$, from equation (12). Our calibration implies that average output is about 6% higher among adjusting firms.




FIGURE 3. Marginal effect of capital reallocation on measured productivity in the steady state.

4. QUANTITATIVE ANALYSIS: AGGREGATE FLUCTUATIONS

This section sets out to answer the questions asked in the Introduction. How much of the cyclical variation in measured productivity is the consequence of capital reallocation? Can endogenous movements in productivity caused by capital reallocation create cyclical variations in measured productivity that resemble business cycles?

One possibility is that capital reallocation is a consequence of variations in aggregate productivity. That is, in response to TFP shocks the gains to reallocation are larger, and hence reallocation is procyclical. A second prospect is the opposite causality: variations in the costs of reallocation lead to changes in measured productivity. In this setting, consistent with the emphasis in Eisefeldt and Rampini (2006), the costs of reallocation are the driving force. Reallocation would be the primary source of variations in aggregate productivity through the induced movements in the Solow residual.

To study these questions, we build on the analysis of the steady state and augment the model with aggregate fluctuations. To the state variables in the planner’s optimization problem, as stated in (3)–(6) we add an aggregate shock—denoted as S , together with its law of motion, detailed below.

We consider two different sources of aggregate fluctuations: (i) TFP shocks and (ii) variations in the cost of reallocation. For both of these cases, we estimate the parameters governing the evolution of the respective aggregate shock. Each shock is assumed to follow an AR(1) process in logs, characterized by an autocorrelation ρ_S , and a standard deviation of the innovations, σ_S . These parameters are chosen to match moments characterizing cyclical capital reallocation. We target (i) the correlation between aggregate productivity and capital reallocation, and (ii) the standard deviation of capital reallocation rates. The resulting parameter estimates are shown in the lower part of Table 1. A comparison between the data moments and the model moments is shown in Table 3.

TABLE 3. Moments: cyclical reallocation.

	Targeted		Untargeted				
	$c(R, \tilde{A})$	$\sigma(R/K)$	$\sigma(\pi)$	$c(\sigma_{\text{ARPK}}, \tilde{A})$	$c(\pi, \tilde{A})$	$\rho_{\tilde{A}}$	$\sigma_{\tilde{A}}$
Data	0.539	0.018	0.027	−0.619	0.368	0.853	0.016
TFP model	0.539	0.019	0.010	−0.840	0.760	0.637	0.015
$\tilde{F}$ model	0.539	0.019	0.011	−1.000	0.940	0.582	0.001

Note: The first two moments are targeted in the calibration. Those include (i) the correlation between the detrended logarithm of measured productivity $\tilde{A}$ and the detrended time series of the logarithm of total capital capital reallocation, (ii) the time-series average of the standard deviation of the fraction of aggregate capital reallocated. The untargeted moments include (i) the time-series average of the standard deviation of the fraction of adjusters, (ii) the correlation between $\tilde{A}$ and the time-series average of the standard deviation of ARPK, (iii) the correlation between $\tilde{A}$ and the fraction of adjusters, (iv) the serial correlation of $\tilde{A}$, and (v) the standard deviation of innovations to $\tilde{A}$.

The first moment in Table 3 indicates that capital reallocation is procyclical, a key fact first emphasized in Eisfeldt and Rampini (2006). Our measure of aggregate productivity is the Solow residual, defined as $\tilde{A}$ in (9). Through the lens of our model, the Solow residual is independent of K and can be constructed from observations on Y and K , using the estimate of α .¹⁹ The second moment in Table 3 is $\sigma(R/K)$. It denotes the standard deviation of the aggregate reallocation rate, a moment which is informative about the magnitude of capital reallocation. Indirectly, by matching this moment we discipline the variability of the shocks that drive capital reallocation. The remaining data moments in Table 3 are not targeted by the model and are discussed below.

4.1 TFP shocks

Consider first the model in which aggregate fluctuations are driven only by shocks to aggregate total factor productivity (TFP). Assume TFP follows an AR(1) in logs:

$$\log A_t = \rho_A \log A_{t-1} + \nu_{A,t}, \quad \nu_A \sim N(0, \sigma_A)$$

TFP shocks are well understood as being able to match a number of business cycle moments. In our setup, they also generate an endogenously procyclical fraction of adjusting firms. During expansions, the gains to reallocation are increasing, while the costs of adjustment are independent of the current value of productivity. This generates procyclical reallocation and amplifies exogenous aggregate fluctuations.

The row labeled “TFP model” in Table 3 shows that this model is able to match these **targeted** reallocation moments, both qualitatively and quantitatively. The cyclicity of reallocation as well as its magnitude are close to their data counterparts.

The mechanism that generates these moments is highlighted by the impulse response functions in Figure 4. A positive shock to TFP (shown in the first panel) has two

¹⁹Our measure of $\tilde{A}$ comoves very closely with alternative empirical measures of aggregate output, such as GDP. All of the data moments reported throughout the paper are robust to using GDP as a measure of the business cycle. We follow Eisfeldt and Rampini (2006) in computing moments with total reallocation, but find that the empirical correlation between total reallocation and reallocation rates (R/K) is 0.89, generating very similar results.




FIGURE 4. Impulse response functions: TFP shock.

effects, a direct effect on measured productivity, and an indirect positive effect on capital reallocation. The magnitude of the resulting increase in reallocation (shown in the bottom-right panel) as well as its positive correlation with the Solow residual (top-right panel) match their empirical counterparts, as shown in the first two columns of Table 3.

The panels in the middle row of Figure 4 highlight the indirect effect of changes in TFP on capital reallocation. More precisely, the panels show the evolution of Γ , the joint distribution of capital and productivity. The blue solid lines show nonadjusting firms, and the red dashed lines show μ_a and ϕ_a that are implied by the optimal choices of the capital vector $\tilde{k}_a$ among adjusting firms. Following an increase in TFP, more capital is assigned to production sites with a high marginal product of capital. This increases ϕ_a and lowers μ_a . Over time, these changes also affect the nonadjusting firms because of time-variation in productivity and adjustment status.

As the bottom-left panel shows, the extensive margin of capital reallocation, denoted as π , is also procyclical (bottom-left panel). A comparison with the data in Table 3 reveals that this procyclicality, measured as $c(\pi, \tilde{A})$, is positive in both the data and the model. However, the model generates a higher correlation of aggregate productivity and the extensive reallocation margin, while underestimating the extensive margin's variability, given as $\sigma(\pi)$ by roughly one-half.

Finally, the model also generates a countercyclical cross-sectional dispersion in firm-level average products of capital, $c(\sigma_{\text{ARPK}}, \tilde{A}) < 0$. Although untargeted by the model, this feature matches the findings of Eisefeldt and Rampini (2006), Kehrig (2015), and others. The model produces this result because following the TFP shock, a larger fraction of firms adjust capital. Adjusting firms have a lower dispersion in average products. Further, the dispersion declines among adjusting and nonadjusting firms because

of changes in the chosen capital vector $\tilde{k}_a$, which are characterized by a higher covariance between capital and productivity ϕ_a .

From Figure 4, the Solow residual, shown in the top-right panel, is not identical to TFP, precisely because of these endogenous dynamics. However, as was anticipated in Section 3.3, these indirect productivity effects are small. We discuss the magnitude further below.

4.2 Variations in the cost of capital adjustment

The second shock we consider is an exogenous variation in the distribution of the cost of capital reallocation. This can drive procyclical reallocation, and thus variations in the Solow residual, while aggregate TFP remains fixed. We assume that $\bar{F}_t$, the upper support of the adjustment cost distribution, follows an AR(1) in logs:

$$\log \bar{F}_t = \log \bar{F} + \rho_{\bar{F}} \log \bar{F}_{t-1} + \nu_{\bar{F},t}, \quad \nu_{\bar{F}} \sim N(0, \sigma_{\bar{F}}).$$

Note that the mean of this process is determined by $\bar{F}$, the value estimated for the steady state of the economy. The parameters $\rho_{\bar{F}}$, and $\sigma_{\bar{F}}$ make up the estimated parameter vector.

The moments generated from simulating the model driven by fluctuations in adjustment costs are shown in the “ $\bar{F}$ model” row of Table 3. The model is able to match both the procyclical reallocation and standard deviation of reallocation moments.

To understand the economic mechanism behind these results, consider a 1% decrease in $\bar{F}_t$. This implies that reallocation becomes cheaper on the margin. The connection between reallocation and productivity stemming from variations in $\bar{F}_t$ is shown in Figure 5.

The decrease in adjustment costs has the immediate effect of leading to an increase in the fraction of adjusting firms, π . This leads to an increase in productivity-enhancing reallocation, which in turn has a positive effect on the Solow residual. In this way, the model is able to generate procyclical capital reallocation and a standard deviation of reallocation rates that matches the data. This mechanism also leads to a countercyclical dispersion in the average products of capital.

Three things distinguish the economy’s response to variations in $\bar{F}_t$ from that to TFP shocks. First, changes in $\bar{F}_t$ cause fully endogenous movements in the Solow residual. Through the dynamics in μ and ϕ induced by lower adjustment costs, the evolution of the joint distribution of productivity and capital Γ generates endogenous movements in the Solow residual. A reduction in $\bar{F}_t$ leads to an increase in π_t , which from (15) increases productivity directly, and from (11) increases the covariance between productivity and capital among nonadjusting firms. In this way, the shock can generate long-lasting effects on measured productivity.

A second distinguishing feature of the $\bar{F}_t$ shock is the shape of the response of measured productivity, which is more detached from the evolution of the exogenous shock than in the TFP case. This response is driven by the dynamics of π and ϕ_n , which as shown in the middle-right panel, respond with a lag. The dynamics highlight the importance of this higher-order moment as an integral component connecting reallocation




FIGURE 5. Impulse response functions: Shock to adjustment costs.

with aggregate productivity. Because changes in $\bar{F}_t$ trigger changes in π_t , this implies that over time, adjusting and nonadjusting firms become more similar (cf. (11)).

Lastly, note that the changes in $\bar{F}_t$ cause comparatively small variations in measured productivity. This is a key reason for why we find that a *quantitatively plausible* degree of variability in $\bar{F}_t$ cannot create the empirically observed cyclicity of the Solow residual, as we show in the following section of the paper.

4.3 Aggregate productivity variation

The previous results showed that both sources of aggregate fluctuations studied were able to match key reallocation moments. However, assessing the contribution of capital reallocation for the variability of the Solow residual, we find that the implications for cyclical variations in the Solow residual coming from the two models are quite different. These differences are highlighted in the last column of Table 3 and in Figure 6. The upshot is that the TFP model matches the standard deviation of the Solow residual almost exactly while the $\bar{F}$ model generates only a fraction of that variation.

From Table 3, the standard deviation of $\tilde{A}$ in the data is 1.6%. In the economy subject to aggregate fluctuations in TFP, calibrated to match the cyclicity of reallocation and the standard deviation of the fraction of reallocated capital, the standard deviation of $\tilde{A}$ is equal to 1.52%.²⁰ Figure 6 decomposes this result and shows that the vast majority of it is the result of variations in TFP.

Here, we uncover a main result of the paper: only about 3% of the standard deviation of $\tilde{A}$ is generated endogenously through capital reallocation. This finding seems

²⁰This moment was not targeted in our calibration of the model.




FIGURE 6. Cyclicalities of aggregate productivity.

to conflict with the steady-state results from Section 3.3, shown in Figure 3. As noted in that discussion, capital reallocation can have potentially large long-run effects on aggregate productivity. But at the estimated parameters that match the U.S. data, we conclude that variations in the cost of capital reallocation cannot create the empirically observed cyclicalities of the Solow residual. Our interpretation of this finding is that the U.S. economy is very efficient in terms of capital reallocation. This reflects both the magnitude of dispersion in productivity across firms as well as the costs of reallocating. Given these estimates, even though reallocation is procyclical, productivity gains from that reallocation are modest.

Importantly, our result does not imply that endogenous changes in capital reallocation can *never* generate sizeable variations in aggregate productivity. Consider a case where adjustment costs significantly increased, leading to a large reduction in the amount capital reallocation. In this case, marginal changes in these adjustments might trigger larger movements in reallocation and productivity. But given the current state of the U.S. economy, reallocation of capital is not a driving force for business cycle dynamics.

One might conjecture that these findings are, in part, a consequence of the specification of capital reallocation costs. Although the specification captures informational frictions in a tractable way, the adjustment behavior is independent of a firm's state. If capital adjustment was fully dependent on the firm state (k, ε) , the average effects of reallocation might even be larger, that is, for a given distribution over (k, ε) . However, with more state dependence in adjustment in the stationary solution, the U.S. economy would appear to be even more efficient. It follows that the marginal effects of changes in the extent of reallocation over the cycle would be even smaller. It is in this sense that our specification of adjustment costs does not drive our results and might, in fact, overstate the effects of procyclical reallocation.

Finally, Figure 3 returns to the economy driven by shocks to adjustment costs and calibrated to match the same data targets. While this economy generates sizeable endogenous movements in the Solow residual, these only amount to 4.5% of the empirical standard deviation in $\tilde{A}$.

4.4 Additional implications

This subsection looks at additional moments from the exercises with aggregate TFP shocks. First, we discuss evidence on the response of reallocation, both on the intensive and extensive margins, to innovations in aggregate TFP. Second, we present moments from an extended version of the model with capital accumulation.

4.4.1 IRFs from the data The next set of results relates TFP shocks to reallocation moments, providing a data counterpart to Figure 4. For this analysis, we estimated local projections as in Jordà (2005), using period- t TFP innovations as exogenous impulses and studied the effects on outcome variations $t + h$ periods ahead. The TFP innovations were taken from Fernald (2014). The estimating equation is

$$y_{t+h} = \gamma_0 + \gamma_h \Delta \text{TFP}_t + \mathbf{d}_h \mathbf{X}_t + \varepsilon_{t+h}, \quad (16)$$

where $h \geq 0$ indicates current and future years. The goal is to estimate, for each horizon h , the sequence of regression coefficients γ_h associated with a time t change in aggregate total factor productivity. The term $\mathbf{X}_t$ denotes control variables, namely lagged values of y_t .²¹ The variable y_t can either be total capital reallocation (in logs) or the fraction of adjusters.

Recall from Figure 4 that, in the model, a TFP innovation led to an immediate increase in both the adjustment rate and the reallocation rate. The adjustment rate then slowly converged back to steady state while the reallocation rate converged back to steady state from below. The data counterparts in Figure 7 are consistent with the impulses: both the adjustment and reallocation rates increase on impact. From the point estimates, both the reallocation rate and the adjustment rate are below steady state after 6 years before rebounding and converging, though not all of this dynamic response is statistically significant.

4.4.2 Adding capital accumulation The baseline model ignores capital accumulation to focus on reallocation. It is of interest to see if the model can also match variations in the intertemporal choice of consumption and investment. There is an important interaction: barriers to reallocation reduce the productivity of accumulating capital, and thus impact the consumption investment margin.

The full extended model with capital accumulation is presented in the Appendix, Section E. In that environment, there is time to build so that investment in period t is

²¹Here, we use the local projections approach to summarize data responses. This was not necessary in determining model responses as those were computed directly from the solution without simulation. We found that varying the number of lags of y_t had no qualitative impact on our results, which are reported with four lags.




FIGURE 7. Local projection estimates.

added to the aggregate capital stock in period $t + 1$. New capital is allocated to existing firms so that the rate of capital accumulation is constant across firms. In this way, accumulation and reallocation are distinct. We solved the capital accumulation model with the baseline parameters, and studied the business cycle moments, thus going beyond the moments characterizing the cyclicity of reallocation. From Table 4, the model with TFP shocks matches both the business cycle features of the data and the frequency of capital reallocation. The estimated serial correlation for the TFP shock and the correlation between output and the Solow residual are higher than in the data but their relative standard deviations are close to the data counterpart. The positive comovement of consumption, investment, output, and the Solow residual is well documented and the model qualitatively reproduces those features.

TABLE 4. Business cycle moments.

	$c(Y, C)$	$c(Y, I)$	$c(C, I)$	$c(Y, \tilde{A})$	$c(Y, \pi)$	$\tilde{\sigma}_C$	$\tilde{\sigma}_I$
Data	0.877	0.714	0.666	0.651	0.338	0.879	3.187
TFP model	0.890	0.857	0.797	0.721	0.407	0.832	1.224

Note: Business cycle data comes from BEA and FRED. All time series are in logs and have been HP-filtered with $\lambda = 100$. The model counterparts come from an economy with aggregate capital accumulation. The symbol $\tilde{\sigma}$ denotes the relative standard deviation of a variable relative to that of aggregate output.

5. CONCLUSION

The goal of this paper was to understand the implications of procyclical capital reallocation for aggregate productivity. The framework was the optimal allocation of a planner facing costs of capital reallocation. Aggregate fluctuations were driven either by TFP shocks or variations in the costs of reallocation.

We have two main conclusions. Our first result is that both a model economy with exogenous productivity shocks and an economy with exogenous variations in the costs of reallocation are able to match key moments regarding the cyclicity of reallocation. Both models create realistic amounts of procyclical capital reallocation, while generating countercyclical productivity dispersion across firms.

However, once the implications for aggregate productivity are taken into account, the two models differ substantially. Our second result is that only the model driven by TFP shocks is able to generate the empirically observed variation in measured aggregate productivity. Importantly, we show that the vast majority of that variation comes from the exogenous shocks themselves. Capital reallocation only contributes by mildly amplifying TFP shocks and does not appear to be playing a major role for U.S. business cycles.

Although the total amount of capital reallocation is an important determinant of long-run productivity, large average levels of reallocation imply small marginal productivity gains. The overall efficiency of capital allocation in the U.S. limits the gains from reallocation over the cycle.

We model capital adjustment frictions as informational costs, allowing us to obtain analytical results. Further, we abstract from other frictions, such as labor reallocation costs, financial constraints, or costs to capital accumulation. Adding those frictions to an environment with state-dependent adjustment behavior would be an interesting avenue for future research.

APPENDIX A: DATA

All data targets used in this paper were computed from Compustat. We use the annual Compustat database between 1971 (the first year our measure of capital reallocation is available) and 2018. We delete non-U.S. firms and firms that only appear during a single year. Further, we drop observations with less than \$100'000 in sales, total assets, total capital, and those with missing total assets or depreciation. We remove firms in the financial industry and utilities (SIC codes 4900–4999 and 6000–6999).²² Firm-level capital is computed using a perpetual inventory method.

For the estimation of revenue-based total factor productivity (TFPR), we follow the methodology of Olley and Pakes (1996) as implemented by İmrohoroğlu and Tüzel (2014). The parameter α governing the curvature of the revenue function in the model is computed as $\alpha = \frac{\hat{\alpha}}{1-\hat{\beta}}$, where $\hat{\alpha}$ and $\hat{\beta}$ denote the estimated elasticities of capital and

²²The main advantage of using Compustat data is that it includes the necessary data items, in particular, the information about capital reallocation. A downside is that it only represents listed firms, which are above average in size, employment, and capital.

labor. These elasticities were estimated controlling for year and 3-digit industry fixed effects. We find $\hat{\alpha} = 0.193$ and $\hat{\beta} = 0.757$, implying $\alpha = 0.81$. This estimate of α is higher than some other estimates. In a study of firm level investment, Cooper and Ejarque (2003) estimated the curvature of the firm revenue function at 0.69. At the firm level, Cooper and Haltiwanger (2006) estimate the curvature at 0.592.

To compute data moments that pertain to capital reallocation, we follow Eisfeldt and Rampini (2006). Capital reallocation R is the sum of sales of property, plant, equipment, and acquisitions. The average fraction of capital being reallocated, R/K , is the average of R divided by total property, plant, and equipment. The average revenue product of capital (ARPK) is defined as Y/K , sales divided by capital. We compute its dispersion as the the time-series median of the annual value-added weighted standard deviations. Investment is the sum of capital expenditures and acquisitions. In the model, we use those same definitions to compare the model to the data.

APPENDIX B: DERIVATIONS

Frictionless economy

Without adjustment costs and with a fixed stock of capital, K , the planner's problem is static:

$$\max_{\tilde{k}_a(\varepsilon)} u(c) \quad (\text{B.1})$$

subject to

$$K = \int_{\varepsilon} \tilde{k}_a(\varepsilon) f(\varepsilon) d\varepsilon, \quad (\text{B.2})$$

$$y = A \int_{\varepsilon} \varepsilon \tilde{k}_a(\varepsilon)^{\alpha} f(\varepsilon) d\varepsilon, \quad (\text{B.3})$$

$$c = y \quad (\text{B.4})$$

The optimal allocation of capital across production sites in the frictionless problem follows from the first-order condition of (B.1), which implies $\alpha A \varepsilon \tilde{k}_a(\varepsilon)^{\alpha-1} = \eta$ for all ε , where η is the multiplier on (B.2). This implies $\tilde{k}_a(\varepsilon) = \frac{\eta}{\alpha A \varepsilon^{\frac{1}{1-\alpha}}}$. Using the constraint (B.2), $\eta = A \alpha K^{\alpha-1} (\int_{\varepsilon} \varepsilon^{\frac{1}{1-\alpha}} f(\varepsilon) d\varepsilon)^{1-\alpha}$. Putting these two conditions together yields the expression for optimal capital in (7).

Stationary distribution

The laws of motion for the elements of the distribution Γ in (10) and (11) imply the stationary values in (12) and (13). The law of motion (11) makes uses of the fact that $\phi' = \text{Cov}(\varepsilon', \tilde{k}_j^{\alpha}) = \text{Cov}(\rho \varepsilon, \tilde{k}_j^{\alpha}) = \rho \phi$ because the covariance between effective capital and productivity among the firms that receive a new random realization of ε' is zero. Using the stationary values (12) and (13) in the Solow residual formulated in (15), we obtain the expression for consumption (14) in the main text. The stationary Solow residual

is given by

$$SR = A \left(\mu_a + \phi_a \cdot \frac{\pi}{1 - \rho_\varepsilon(1 - \pi)} \right). \quad (\text{B.5})$$

APPENDIX C: SOLUTION METHOD

The planner's problem consists of maximizing the present-discounted stream of utility $\sum_{t=0}^{\infty} \beta^t \log(c_t)$. The choice variables are a capital vector $\tilde{k}_a$ for firms whose capital stock is adjusted, as well as the fraction of adjusters π .

Steady state

To solve the steady state of the model with adjustment costs, the choice variables are constants $\pi_t = \pi$ and $\tilde{k}_{a,t} = \tilde{k}_a \forall t$. The associated moments of the capital vector are μ_a and ϕ_a , with production in adjusting firms given by $y^a = \mu_a + \phi_a$. The steady-state values of the moments of the nonadjusters are given by (12) and (13). Aggregate productivity A is constant and normalized to one. The steady state of the model is characterized by a constant consumption level $c = \mu_a + \phi_a \frac{\pi}{1 - \rho_\varepsilon(1 - \pi)}$, as laid out in the main text.

By iteratively substituting out future state variables μ_n and ϕ_n with their laws of motion and initial values μ_0 and ϕ_0 , one can trace out the dynamic impact of all choice variables on production in all future periods. For example, the choice of π is dynamic because it affects the evolution of the state variables μ_n and ϕ_n in addition to affecting the output mix between adjusters and nonadjusters as well as the adjustment costs.

To find the steady-state value of π , we take the first-order condition of the present discounted stream of log consumption with respect to π . It is given by

$$\frac{1}{c(\pi)} \cdot \left[A \phi_a \frac{1 - \rho_\varepsilon}{1 - \beta \rho_\varepsilon(1 - \pi)} - \pi \bar{F} \right] = 0. \quad (\text{C.1})$$

Here, the dependence of consumption on π is made explicit. The term in the denominator results from the infinite sum. For the steady state, this is a cubic equation, which has the real solution

$$\pi = \frac{1}{3} \cdot a_2 + S + T, \quad (\text{C.2})$$

with parameter values given by $a_2 = (1 - \beta(2 \cdot \rho_\varepsilon - 1))/(\rho_\varepsilon \beta)$, $S = \sqrt[3]{R + \sqrt{D}}$, $T = \sqrt[3]{R - \sqrt{D}}$, $R = (9a_1 a_2 - 27a_0 - 2a_2^3)/54$, $a_1 = (1 - \rho_\varepsilon(1 + \beta) + \beta \rho_\varepsilon^2)/(\rho_\varepsilon^2 \beta)$, $a_0 = -\phi_a(1 - \rho_\varepsilon)/(\bar{F} \rho_\varepsilon^2 \beta)$, and $D = ((3a_1 - a_2^2)/9)^3 + ((9a_1 a_2 - 27a_0 - 2a_2^3)/54)^2$.

Similarly, we can solve for $\tilde{k}_a$ by taking the derivative of the presented-discounted stream of log consumption. In addition, there is a Lagrange multiplier λ , which denotes the constraint that the sum of capital must always equal K , normalized to one. The first-order condition for capital at firm i is given by $k_i^{\alpha-1} \zeta_i = \lambda$, with

$$\zeta_i = \frac{A \pi \alpha}{c(\tilde{k}_a)} \cdot \left[\frac{\varepsilon_i - 1}{1 - \beta \rho_\varepsilon(1 - \pi)} + \frac{1}{1 - \beta(1 - \pi)} \right] \quad (\text{C.3})$$

This allows us to write the multiplier as

$$\lambda = \left(\int_i \xi_i^{\frac{1}{1-\alpha}} d\Gamma(k, \varepsilon) \right)^{1-\alpha} \quad (\text{C.4})$$

and solve for the optimal capital vector.

Aggregate shocks

To solve for the economy with aggregate shocks, we make use of the fact that the state of our economy can be written in terms of current choice variables and the initial states of μ_n and ϕ_n . This allows us to explicitly compute the effect of current choices on future values without relying on approximations or grids for the endogenous variables. For example, to find the optimal level of π in period $t = 0$, we can rewrite consumption in period t as follows. For this illustration only, we have made the dependence on variables at time 0 and time 1 explicit and assumed that starting in $t = 2$ all choice variables are at their steady-state values. This is not the case in the solution of the aggregate model, where we assume that the model returns to back to steady state in period T , where T is chosen so as to have no effect on the solution:

$$\begin{aligned} c_t = & \mu_a \pi (1 + (1 - \pi) + \cdots + (1 - \pi)^{t-2}) + \phi_a \pi (1 + (1 - \pi) \rho_\varepsilon + \cdots + (1 - \pi)^{t-2} \rho_\varepsilon^{t-2}) \\ & + (1 - \pi)^{t-1} \pi_1 (\mu_{a1} + \phi_{a1} \rho^{t-1}) + (1 - \pi)^{t-1} (1 - \pi_1) \pi_0 (\mu_{a0} + \phi_{a0} \rho^{t-1}) \\ & + (1 - \pi)^{t-1} (1 - \pi_1) (1 - \pi_0) (\mu_0 + \phi_0 \rho^{t-1}) - \frac{\pi^2}{2} \bar{F} \end{aligned} \quad (\text{C.5})$$

One can easily verify that this expression corresponds to what was derived for the steady state in the main text by letting $t \rightarrow \infty$, which implies that production converges to $\mu_a + \phi_a \frac{\pi}{1 - \rho_\varepsilon(1 - \pi)}$, as stated. The first-order condition for π_0 is given by

$$(\mu_{a0} - \mu_0) \cdot \frac{1 - \beta(\pi_1 - \pi)}{1 - \beta(1 - \pi)} + (\phi_{a0} - \phi_0) \cdot \frac{1 - \beta\rho(\pi_1 - \pi)}{1 - \beta\rho(1 - \pi)} - \pi_0 \cdot \bar{F} = 0. \quad (\text{C.6})$$

To find the optimal capital assignment vector, one can rewrite the present discounted stream of consumption in a similar manner. The idea is, once more, that as all future choices are known, the effect of current choices on future endogenous states can be determined exactly. The optimal choices of π and capital have to be consistent. In the computational implementation, we therefore begin with a guess for π and then solve for the optimal capital vector. This implies a value for the mean and covariance terms μ_a and ϕ_a , which are used in the solution for π in (C.6). The updated value for π is then used in a new solution for the capital vector until convergence.

APPENDIX D: SENSITIVITY ANALYSIS

The SMM estimation of the baseline parameters resulted in an extremely close fit between the model and the targeted and untargeted data moments (cf. Table 2). In this

TABLE D.1. Sensitivity analysis.

Parameter	$R > 0$	R/K	σ_{TFPR}
β	0.1183	0.0005	−0.0019
α	1.3139	0.1413	−0.9621
σ_ε	0.9933	0.0676	0.3485
$\bar{F}$	−1.3577	−0.0168	0.10417
ρ_ε	−2.997	−0.8193	−0.0318

Note: The first column shows the parameter that is changed. All other parameters are held fixed at their estimated baseline values. The remaining columns show the numerical derivatives of the three targeted moments with respect to a change in a parameter.

section, we conduct a number of robustness exercises that show the impact of key exogenously set parameters on the fit of the model and on our main conclusions.

An important parameter is α , which determines the span-of-control in our model. A lower value α implies a larger gain from spreading capital across production sites. We reestimate our model with a value of $\alpha = 0.7$ commonly used in the literature. The fit of the targeted moments is virtually unchanged.²³ In terms of the untargeted moments, the model with a lower α generates very comparable numbers, with the serial correlation of firm-level sales being slightly closer to the data than the baseline model. Most notably, the fraction of small capital adjustment rates decreases from 86.8% in the baseline, to 68.1%. Our conclusions are hardly affected by the lower α . The steady-state results still imply sizeable potential long-run losses from reductions in capital reallocation if the fraction of adjusters were exogenously decreased to 25% of its currently observed levels (7.4% vs. 12.7% in the baseline). The results with aggregate fluctuations still imply that while the model with TFP shocks generates virtually the same standard deviation in measured aggregate productivity as what is observed in the data, the model with shocks to the costs of adjustments only accounts for about 4.6% of that variation.

More generally, Table D.1 shows how sensitive the model moments are with respect to changes in the parameter vector. The table shows numerical partial derivatives of the model moments with respect to the parameter listed in the first column.

APPENDIX E: CAPITAL ACCUMULATION

This section extends the baseline model to include capital accumulation. The planner’s problem in (3) is modified in the following way. Given the state, the planner makes an investment decision K' , determines π , and chooses how to allocate capital among adjusting firms, $\tilde{k}_a(k, \varepsilon)$. The notation is unchanged with respect to the main text. The choice problem of the planner is

$$V(\Gamma(k, \varepsilon), K) = \max_{\pi, \tilde{k}_a(k, \varepsilon), K'} u(c) + \beta E_{\Gamma'|\Gamma} V(\Gamma'(k, \varepsilon), K')$$

(E.1)

²³The resulting parameter estimates were $\bar{F} = 0.102$, $\rho_\varepsilon = 0.896$, and $\sigma_\varepsilon = 0.373$. For the economy with aggregate fluctuations, we find $\rho_A = 0.702$, $\sigma_A = 0.017$, $\rho_{\bar{F}} =$, $\sigma_{\bar{F}} =$. The benchmark values are reported in Table 1.

subject to the constraint that capital in all production sites sum up to K , given by (5), and the resource constraint (6), amended to include investment:

$$c + K' = y + (1 - \delta)K - K \int_0^{F(\pi)} F dG(F), \quad (\text{E.2})$$

with output y defined in (4). Because capital is firm specific, it is necessary to specify transition equations at the firm level. We model capital accumulation in a manner intended to distinguish reallocation from aggregate capital accumulation. Hence, we assume that the capital at all firms, regardless of their reallocation status, have the same capital accumulation. In this way, all firm specific capital reflects reallocation rather than firm specific accumulation. Put differently, the model decomposes the accumulation of capital at a particular firm into two components: (i) the aggregate investment rate and (ii) the reallocation of capital. If there were no adjustment costs, this decomposition would be innocuous. With adjustment costs associated with information about a firm's current state, the costs of reallocation are prominent.

Specifically, let $i = \frac{K' - K(1-\delta)}{K}$ denote the gross investment rate so that $K' = (1 - \delta + i)K$ is the aggregate capital accumulation equation. The transition for capital this period (after reallocation) and the initial firm-specific capital next period is given by

$$k'_j(k, \varepsilon) = (1 - \delta + i)\tilde{k}_j(k, \varepsilon), \quad (\text{E.3})$$

for $j = a, n$.

REFERENCES

- Bachmann, Rüdiger and Christian Bayer (2014), "Investment dispersion and the business cycle." *American Economic Review*, 104, 1392–1416. <https://www.aeaweb.org/articles?id=10.1257/aer.104.4.1392>. [1342]
- Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta (2013), "Cross-country differences in productivity: The role of allocation and selection." *American Economic Review*, 103, 305–334. <https://www.aeaweb.org/articles?id=10.1257/aer.103.1.305>. [1342]
- Bloom, Nicholas (2009), "The impact of uncertainty shocks." *Econometrica*, 77, 623–685. <https://doi.org/10.3982/ECTA6248>. [1342]
- Caballero, Ricardo J. and Eduardo M. R. A. Engel (1999), "Explaining investment dynamics in u.s. manufacturing: A generalized (s, s) approach." *Econometrica*, 67, 783–826. <https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0262.00053>. [1350]
- Chen, Kaiji and Zheng Song (2013), "Financial frictions on capital allocation: A transmission mechanism of tfp fluctuations." *Journal of Monetary Economics*, 60, 683–703. [1342]
- Cooper, Russell and João Ejarque (2003), "Financial frictions and investment: Requiem in q." *Review of Economic Dynamics*, 6, 710–728. <http://www.sciencedirect.com/science/article/pii/S1094202503000644>. [1361]

Cooper, Russell W. and John C. Haltiwanger (2006), “On the nature of capital adjustment costs.” *The Review of Economic Studies*, 73, 611–633. <https://doi.org/10.1111/j.1467-937X.2006.00389.x>. [1344, 1345, 1361]

Cooper, Russell W., and Immo Schott (2023), “Supplement to ‘Capital reallocation and the cyclical nature of aggregate productivity.’” *Quantitative Economics Supplemental Material*, 14, <https://doi.org/10.3982/QE1892>. [1339]

Eisfeldt, Andrea L. and Adriano A. Rampini (2006), “Capital reallocation and liquidity.” *Journal of Monetary Economics*, 53, 369–399. <http://www.sciencedirect.com/science/article/pii/S0304393206000298>. [1339, 1340, 1342, 1344, 1352, 1353, 1354, 1361]

Fernald, John G. (2014), “A quarterly, utilization-adjusted series on total factor productivity.” Report, Federal Reserve Bank of San Francisco <https://ideas.repec.org/p/fip/fedfwp/2012-19.html>. [1358]

Foster, Lucia, John Haltiwanger, and Chad Syverson (2008), “Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?” *American Economic Review*, 98, 394–425. <https://www.aeaweb.org/articles?id=10.1257/aer.98.1.394>. [1339, 1341]

Guner, Nezih, Gustavo Ventura, and Yi Xu (2008), “Macroeconomic implications of size-dependent policies.” *Review of Economic Dynamics*, 11, 721–744. <http://www.sciencedirect.com/science/article/pii/S1094202508000070>. [1341]

Hsieh, Chang-Tai and Peter J. Klenow (2009), “Misallocation and manufacturing tfp in China and India.” *The Quarterly Journal of Economics*, 124, 1403–1448. <https://doi.org/10.1162/qjec.2009.124.4.1403>. [1339, 1341]

İmrohoroglu, Ayşe and Şelale Tüzel (2014), “Firm-level productivity, risk, and return.” *Management Science*, 60, 2073–2090. <https://doi.org/10.1287/mnsc.2013.1852>. [1360]

Jordà, Òscar (2005), “Estimation and inference of impulse responses by local projections.” *American Economic Review*, 95, 161–182. <https://www.aeaweb.org/articles?id=10.1257/0002828053828518>. [1358]

Kaymak, Barış and Immo Schott (2019), “Loss-offset provisions in the corporate tax code and misallocation of capital.” *Journal of Monetary Economics*, 105, 1–20. <http://www.sciencedirect.com/science/article/pii/S0304393219300777>. [1341]

Kehrig, Matthias (2015), “The cyclical nature of the productivity distribution.” Report, SSRN Working paper <https://ideas.repec.org/p/cen/wpaper/11-15.html>. [1342, 1354]

Klenow, Peter J. and Jonathan L. Willis (2007), “Sticky information and sticky prices.” *Journal of Monetary Economics*, 54, 79–99. [1344]

Krusell, Per and Anthony A. Jr. Smith (1998), “Income and wealth heterogeneity in the macroeconomy.” *Journal of Political Economy*, 106, 867–896. www.jstor.org/stable/10.1086/250034. [1342, 1343, 1347]

Lanteri, Andrea (2018), “The market for used capital: Endogenous irreversibility and reallocation over the business cycle.” *American Economic Review*, 108, 2383–2419. <https://www.aeaweb.org/articles?id=10.1257/aer.20160131>. [1342, 1348]

Lucas, Robert E. (1978), "On the size distribution of business firms." *The Bell Journal of Economics*, 9, 508–523. <http://www.jstor.org/stable/3003596>. [1343]

Maksimovic, Vojislav and Gordon Phillips (2001), "The market for corporate assets: Who engages in mergers and asset sales and are there efficiency gains?" *The Journal of Finance*, 56, 2019–2065. <https://onlinelibrary.wiley.com/doi/abs/10.1111/0022-1082.00398>. [1339, 1341]

Olley, G. Steven and Ariel Pakes (1996), "The dynamics of productivity in the telecommunications equipment industry." *Econometrica*, 64, 1263–1297. <http://www.jstor.org/stable/2171831>. [1339, 1341, 1346, 1349, 1360]

Osootimehin, Sophie (2019), "Aggregate productivity and the allocation of resources over the business cycle." *Review of Economic Dynamics*, 32, 180–205. <http://www.sciencedirect.com/science/article/pii/S1094202518302709>. [1341]

Ramey, Valerie A. and Matthew D. Shapiro (2001), "Displaced capital: A study of aerospace plant closings." *Journal of Political Economy*, 109, 958–992. <http://www.jstor.org/stable/10.1086/322828>. [1341]

Restuccia, Diego and Richard Rogerson (2008), "Policy distortions and aggregate productivity with heterogeneous establishments." *Review of Economic Dynamics*, 11, 707–720. <http://www.sciencedirect.com/science/article/pii/S1094202508000203>. [1339, 1341]

Thomas, Julia K. (2002), "Is lumpy investment relevant for the business cycle?" *Journal of Political Economy*, 110, 508–534. <https://www.journals.uchicago.edu/doi/abs/10.1086/339746>. [1348]

Co-editor Kjetil Storesletten handled this manuscript.

Manuscript received 7 April, 2021; final version accepted 15 May, 2023; available online 24 May, 2023.