Aggregate Dynamics and Microeconomic Heterogeneity: The Role of Investment-Specific Technology*

Giuseppe Fiori Filippo Scoccianti

March 19, 2025

Abstract

We explore how firm heterogeneity affects business cycle dynamics within a general equilibrium model when technological progress is investment-specific. Our data show that investment age—the time elapsed since the firm experienced a significant investment episode—reflects technology obsolescence and is negatively correlated with firms' total factor productivity. In the model, new technologies are embodied in new capital goods, and non-convex adoption costs limit firms' access to new technologies. Idiosyncratic and aggregate shocks alter the technology adoption timing, leading to procyclical movements in aggregate productivity that amplify aggregate dynamics compared to a neoclassical growth model.

JEL Codes: D24; E22; E32.
Keywords: Business Cycles; (S,s) policies; Vintage Effects; Firm Heterogeneity; Lumpy investment.

*Giuseppe Fiori: Board of Governors of the Federal Reserve System, Division of International Finance, 20th and C St. NW, Washington D.C. 20551, United States (email: giuseppe.fiori@frb.gov; webpage: http://www.giuseppefiori.net); Filippo Scoccianti: Banca d'Italia, Via Nazionale 91, Rome, Italy (email: filippo.scoccianti@bancaditalia.it). We are grateful to Andrea Caggese, Federico Cingano, Marcela Eslava, Jesús Fernández-Villaverde, Domenico Ferraro, Reint Gropp, Michael Koetter, Massimo Giovannini, David Jenkis, Aubhik Khan, Don Koh, Andrea Lanteri, Francesca Lotti, Francesco Manaresi, Plutarchos Sakellaris, Jasmine Santana, Paolo Sestito, Fabio Schiantarelli, Julia Thomas, and seminar and conference participants at the Board of Governors, Macroeconomic Dynamics Conference, North Carolina State University, TDM Workshop at Duke University, Midwest Macro Meetings, the Ohio State University, the University of Arkansas, and FINPRO-CompNet Workshop for useful comments. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System (or of any other person associated with the Federal Reserve System) or the Bank of Italy.

1 Introduction

Macroeconomic models that emphasize firm heterogeneity typically assume that productivity differences across firms are random; see, for instance, Hopenhayn (1992). Since the work of Johansen (1959) and Solow (1960), a large body of the theoretical literature has emphasized capital-embodied technological progress, because newer vintages of investment goods are of better quality and may enhance the efficiency of existing capital.

In this paper, we examine the role of capital-embodied technological progress for firms' productivity heterogeneity and business cycle dynamics. We document that investment age—the time elapsed since a firm experienced a significant investment episode—proxies for the vintage technology operated by the firm. In turn, investment age is negatively correlated with a firm's total factor productivity (TFP). We then study the aggregate implications of heterogeneity in investment age and formulate a state-of-the-art model of lumpy investment that builds on Khan and Thomas (2003, 2008) in which a non-convex adoption cost prevents firms from switching from less efficient to newer, more efficient, investment goods. As firms invest in different technologies, the equilibrium features a non-degenerate distribution of capital stocks and technologies across firms. Idiosyncratic and aggregate shocks alter the timing of investment in the latest technologies at the firm level and result in shifts in the cross-sectional distribution of capital and investment-specific productivity. This shift generates endogenous procyclical movements in economy-wide productivity of investment goods beyond the ones accounted for by exogenous shocks to the efficiency of the latest vintage. The endogenous shift in productivity implies that microeconomic heterogeneity in investment age *amplifies* the magnitude of business cycle fluctuations relative to a benchmark neoclassical growth model. Unlike in work by Thomas (2002) and Khan and Thomas (2008), in our framework, lumpy investment, resulting in a technology adoption decision, enhances the response of aggregate variables to technology shocks beyond the initial impulse of the shock.

We start by documenting the nature of capital accumulation at the firm level and its importance for aggregate investment dynamics. Using 30 years of data that cover over two-thirds of the value added in the Italian economy, we show that investment at the firm level is a large and infrequent, or *lumpy*, episode—in line with many studies across advanced economies.[1] On average, every year only 20 percent of firms exhibit investment spikes or an investment rate above 20 percent, but they account for about two-thirds of total investment in our data.

Using data from the Survey of Industrial and Service Firms (INVIND) data, we show that firms with higher investment age self assess the technology they operate as more obsolete relative to the technological frontier. Technology obsolescence is also displayed in measured TFP, with firms' investment age negatively correlated with Solow residuals. This negative relationship holds accounting for time, industry, time-industry, and firm-specific effects and managers' expectations about future sales. The vast heterogeneity in investment age in the data points to widespread differences in technologies available to firms, with the average firm displaying an investment age equal to about three and a half years.

To study the role of investment age heterogeneity for aggregate dynamics, we formulate a general equilibrium framework that builds on Khan and Thomas (2003, 2008) where new technologies are embodied in new investment goods and firms are subjected to exogenous idiosyncratic shocks. The distinctive feature of the model is a quality ladder structure in investment-specific productivity. In the model, the firm's productivity, which maps physical units of investment on to capital services, is endogenous to the timing of technology adoption. Every period, firms optimally decide when to invest in the latest vintage of capi-

[1] The coverage in terms of value added has been increasing over time, from around 60 percent at the beginning of the '90s to around 80 percent at the end of the sample period.

tal goods and, therefore, the associated latest technology. As this choice is subject to a non-convex adoption cost, the firm's policy functions are of the (S,s) type: Some firms buy the newest capital goods and adopt the latest technology, while others postpone it and keep buying the current vintage they operate. Conditional on the adoption decision, firms optimally choose investment. In equilibrium, technologies of different quality coexist. The history of adoption choices contributes to microeconomic heterogeneity in productivity and, in turn, determines aggregate productivity. Instead, when the adoption cost is set equal to zero, all the firms find it optimal to adopt the latest technology in every period, with our framework boiling down to a standard neoclassical growth model with investment-specific technical progress as in Greenwood, Hercowitz and Krusell (1997, 2000), where firms' and aggregate productivity are exogenous. What differentiates our model from Khan and Thomas (2003) is the ability of firms to invest in capital goods embodying different efficiencies. In Khan and Thomas (2003) and in standard models of investment-specific technical progress such as Solow (1960) and Greenwood, Hercowitz and Krusell (2000), firms have the option to invest only in capital goods embodying the latest vintage.

Our framework accounts well for microeconomic heterogeneity in investment by reproducing the cross-sectional distributions of investment rates and investment age and the near-zero autocorrelation of investment rates. Also, the model accounts for about 10 percent of the cross-sectional dispersion in idiosyncratic TFP measured in the data, with the vintage technology structure contributing to increased dispersion beyond the one implied by the exogenous component of TFP.

The main result of the paper is that microeconomic heterogeneity is a key determinant of business cycle dynamics following investment-specific productivity shocks. In response to an exogenous increase in the efficiency of the frontier, the canonical investment-specific shock, firms increase investment by acquiring capital embodying a more efficient technology. The ensuing increase in aggregate productivity is larger than the one implied by the shock as firms switch from less efficient capital goods to more efficient ones. From a quantitative standpoint, this channel boosts aggregate productivity, amplifying aggregate series' persistence and volatility relative to a standard neoclassical model with investment-specific technological progress. Our findings contrast with the earlier work by Thomas (2002) and Khan and Thomas (2003, 2008), where including firm heterogeneity is irrelevant for aggregate dynamics relative to a neoclassical benchmark because of general equilibrium effects that dampened the increase in investment. In our framework, general equilibrium effects do not neutralize firm heterogeneity because of the productivity gains associated with the investment decision. As the latest technology is more (less) efficient than the previous one, the current technology available to the firm becomes more (less) obsolete relative to the technological frontier, increasing (decreasing) the benefit of adoption. As more firms adopt the latest vintage, shifts in the distribution of capital stocks and technologies across firms induced by the aggregate shock lead to fluctuations in the economy-wide productivity. The endogenous response of productivity constitutes an additional force that amplifies fluctuations in investment and macroeconomic dynamics relative to a neoclassical growth model. For a given persistence, the vintage technology model requires shocks about 60 percent smaller than the neoclassical growth model to generate business cycles of the same magnitude. Our results support the view that microeconomic heterogeneity is relevant for aggregate dynamics as investment contributes to productivity.

Our paper is organized as follows. In Section 1.1, we describe our contribution relative to the existing literature. We document the nature of capital accumulation at the firm level in Section 2 and the relationship between capital accumulation and productivity in Section 3. Section 4 outlines the model incorporating vintage technology and rich firm heterogeneity. In Sections 5 and 6, we describe the parameterization of the model and its quantitative performance relative to the data. In Section 7, we quantify the role of heterogeneity in vintage technology for aggregate dynamics. Section 8 concludes.

1.1 Literature Review

Our work connects to different strands of the empirical and theoretical literature on capital-embodied technological progress, firm heterogeneity, and business cycle dynamics. Our analysis documents the link between capital accumulation and productivity using firm-level data. After Gordon (1990) and Cummins and Violante (2002), who use product-level and sectoral data, most of the existing literature on vintage capital has focused on aggregate data; see, for instance, Hulten (1992); Wolff (1996); and Greenwood, Hercowitz and Krusell (1997, 2000). The central insight of these papers is that, under some conditions, the growth rate of the price of investment goods (relative to the one of consumption) is a measure of investment-specific technological progress.[2] There is, however, little systematic evidence on the role of capital accumulation for productivity dynamics at the *firm level* partly because a rigorous analysis requires data not commonly available to researchers. Licandro, Maroto Illera and Puch (2005), Power (1998), and Sakellaris and Wilson (2004) are the exceptions.[3] We extend the existing studies by showing that investment age is a proxy of technology at the firm level. We explicitly connect investment age to TFP heterogeneity to analyze its aggregate implications.

Our focus on the business cycle implications of microeconomic heterogeneity relates our work to the literature that studies sectoral and aggregate dynamics in models with rich firm heterogeneity; see, for instance, Cooper and Haltiwanger (1993), Caballero and Engel (1999),

[2] Knittel (2011) and Bertolotti, Gavazza and Lanteri (2023) estimate the evolution of the technological frontier in cars without using prices and by positing a marginal-cost function that depends on vehicle attributes.

[3] Using U.S. manufacturing data, Power (1998) finds no evidence that investment spikes contribute to increasing a firm's productivity. Using similar data Sakellaris (2004), Sakellaris and Wilson (2004), and, more recently (and using Spanish manufacturing data) Licandro, Maroto Illera and Puch (2005) find the opposite result.

Khan and Thomas (2008), and Bachmann, Caballero and Engel (2013), to name a few. We retain several elements that have determined the quantitative success of this class of models in accounting for the pattern of capital accumulation at the firm level. Also, in our model, at the firm level, TFP is endogenous as firms decide when to invest in the latest vintage of capital goods. Thus, the vintage technology increases TFP dispersion beyond the one implied by exogenous productivity.

The literature has also debated the relevance of accounting for the cross-sectional dynamics in investment for aggregate dynamics; see Thomas (2002), Khan and Thomas (2008), Fiori (2012), Bachmann, Caballero and Engel (2013), House (2014), and Winberry (2021). We contribute to this literature by showing that technology adoption motives constitute a powerful amplification mechanism of aggregate disturbances, providing a channel through which microeconomic heterogeneity shapes business cycle dynamics.

Our model connects with the literature focused on business cycle dynamics. In estimated dynamic stochastic general equilibrium (DSGE) models, as in Smets and Wouters (2007) and Justiniano, Primiceri and Tambalotti (2010), technology shocks are typically important drivers of aggregate dynamics at the business cycle frequency. We show that firm heterogeneity is a quantitatively-relevant feature of the propagation of technology innovation. On a similar note, we find that the procyclicality of adoption also generates a countercyclical dispersion in TFP without relying on the stochastic volatility of productivity as in Bachmann and Bayer (2014) and Bloom et al. (2018).

The main difference from models that study vintage capital based on Solow (1960) is that, in our framework, investment goods embodying different efficiencies available to firms. The quantitative focus of our analysis distinguishes our work from the abundant literature on vintage capital. As in Solow (1960), in general, the dynamics of capital vintage models cannot be captured through a representative firm unless knife-edge conditions are met—for instance, constant returns to scale in production. As a result, the number of studies that have confronted vintage models with microeconomic data has been limited. For a complete list of references and a historical perspective on the evolution of the literature on vintage capital, see the extensive survey of Boucekkine, de la Croix and Licandro (2011).

2 Firm Heterogeneity in Investment Age and Technology

This section documents the relationship between capital accumulation and vintage technology at the firm level. After discussing our data sources in Section 2.1, we show in Section 2.2 that capital accumulation at the firm level is a large and infrequent, or lumpy, episode. In Section 2.3, using survey data, we provide evidence that the time elapsed between lumpy episodes of investment, or investment age, proxies for how obsolete the technology operated by firms is relative to the technological frontier. In Section 2.4, we construct the cross-sectional distribution of investment age and show that it points to widespread differences in vintage technology across firms.

2.1 Description of Data Sources

We obtained our data set by combining different sources. To construct the main variables of interest, firm-level investment rates and measures of productivity, we require information on payroll, gross value-added, and employment taken from yearly balance sheets from the Cerved Group S.P.A. (Cerved Database), INVIND, and the Italian National Institute for Social Security (INPS) (see Appendixes A and B for detailed information on data sources and variables construction).

Cerved contains firms' balance sheet information. The database spans 30 years, from 1986 to 2016, and matches the size and the distribution of Italian firms accounting for up to

80 percent of the value added produced in the Italian economy. Consistent with their share of the economy, the manufacturing and trade sectors constitute more than one-half of the observations in the data.[4]

To complement the data in Cerved, we employ INVIND an annual business survey that elicits firms' expectations and contains information on investment and available technology. INVIND is conducted between February and April of every year by the Bank of Italy on a representative sample of firms operating in industrial sectors (manufacturing, energy, and extractive industries), construction, and nonfinancial private services, with the administrative headquarters in Italy.[5] The available sample extends from 1994 to 2016. INVIND contains information about firms' self-assessments of the technology they operate that allows us to validate investment age as a proxy for technology. Also, INVIND includes information on managers' expectations about future sales that we exploit to control for potential confounding factors related to unobservable time-varying effects driving the relationship between investment age and TFP.

2.2 The Lumpy Nature of Capital Accumulation

We now document the lumpy nature of capital accumulation at the firm level by computing the distribution of investment rates for the Cerved database. As is customary in the literature, we calculate real capital stocks by applying a perpetual inventory method from balance sheet data (see Appendix B for details). Following Bloom (2009), we define the investment rate for a given firm f at time t as $ik_{f,t} = \frac{I_{f,t}}{0.5(K_{f,t-1}+K_{f,t})}$, where $I_{f,t}$ is the real investment net of disinvestment. Investment $I_{f,t}$ includes expenditures on equipment and structures as they

[4] In Table A.1, we report the composition of the data set by sector. The sectors are identified following the statistical classification of economic activities in the European Community, abbreviated as NACE.

[5] Specifically, INVIND represents the Italian economy based on the branch of activity (according to an 11-sector classification), size class, and region in which the firm's head office is located.

are not separately identifiable in our data. Table 1 reports the empirical distribution of $ik_{f,t}$ in our Cerved sample (statistics based on the smaller matched INVIND sample are similar). As in Bachmann and Bayer (2014), among others, we define lumpy adjusters as those firms

Table 1: Cross-Sectional Distribution of Investment Rates

Investment Rate	Share in Data Set (A)	Share of Investment (B)	Share of Output (C)	Share of Employment (D)
$ik \geq 20\%$	20.04%	64.01%	26.77%	27.52%
$ik \geq 5\%$	57.99%	101.62%	70.51%	68.91%
$-5\% \leq ik \leq 5\%$	37.21%	5.76%	25.67%	27.01%
$ik \leq -5\%$	4.80%	-8.02%	3.82%	4.08%
$ik \leq -20\%$	2.32%	-6.65%	1.98%	2.14%

Note: ik denotes the investment rate. See the main text for the definition. The distribution of investment rates is computed using the Cerved database over the sample period 1998 to 2016.

that exhibit a positive spike—i.e., an investment rate above 20 percent. On average, these investors account for about two thirds of total investment. Firms that experience small capital adjustments (defined as in Øivind and Schiantarelli (2003) as experiencing $ik_{f,t}$ between negative 5 and 5 percent) account for only 6 percent of total investment.[6] When grouping firms by their investment rates, the share of investment that the individual groups of firms account for differs substantially; instead, the shares of output and employment are similar.

[6] The lumpy nature of the capital accumulation process is also a feature of the data in other countries. Doms and Dunne (1998) report evidence for the United States; Bachmann and Bayer (2014) for Germany; Licandro, Maroto Illera and Puch (2005) for Spain; Øivind and Schiantarelli (2003) for Norway; and Gourio and Kashyap (2007) for Chile.

2.3 Investment Age As a Proxy for Firm's Technology

In this subsection, we construct a firm-level measure of investment age based on the timing of investment spikes. We show that investment age proxies the technology available to the firm, supporting the idea that part of the technological progress is embodied in new capital.

Investment age, denoted by $Inv.Age_{f,t}$, is based on the time elapsed since the last *positive* investment spike experienced by each firm. As discussed in the previous subsection, we define an investment spike using a threshold of 20 percent. When a firm experiences a positive investment spike, the variable $Inv.Age_{f,t}$ equals zero, progressively increasing by one each year until the same firm experiences an investment spike. The long time-series dimension in our data allows us to split the sample using roughly 40 percent to initialize the distribution of $Inv.Age_{f,t}$ and the remaining part of the sample in the empirical analysis. As Cerved and INVIND span different dates, the sample periods considered are not aligned.

Investment Age and Vintage Technology We employ INVIND survey data to validate the interpretation of $Inv.Age$ as a measure of vintage technology—i.e., how distant the technology available to the firm is from the frontier. Specifically, the 2014 wave of INVIND asked surveyed firms to self-assess how advanced their technology was using a discrete index from 4, indicating the availability of the most advanced technology, to 1, indicating obsolete technology.

Using the 2014 cross-section, we regress the index of the firm's technology $Tech.Adv_{f,2014}$ on investment age and estimate a negative correlation: Firms with higher investment age employ less advanced technologies. Specifically, $Tech.Adv_{f,2014} = 3.529 - 0.014 \times Inv.Age_{f,2014} + \epsilon_{f,2014}$, where the coefficient on investment age is statistically significant at 5 percent (p-value equal to 0.025) and the constant is significant at 1 percent.[7] This cross-sectional evi-

[7]The regression is estimated using the matched sample Cerved-INVIND for 2014 that responded to the question of 181 observations.

dence supports the interpretation of $Inv.Age_{f,t}$ as a firm-level measure of the distance from the technological frontier.

2.4 Investment Age Heterogeneity

The cross-sectional distribution of investment age in the data, obtained by averaging across the sample period, points to substantial heterogeneity in firms' investment age and indicates widespread differences in technologies available to firms (see Figure 1). As described in the previous section, the fraction of firms that display an investment age equal to zero has experienced an investment spike in the current period. As firms delay a significant capital adjustment, investment age progressively increases.

Figure 1: Investment Age Distribution

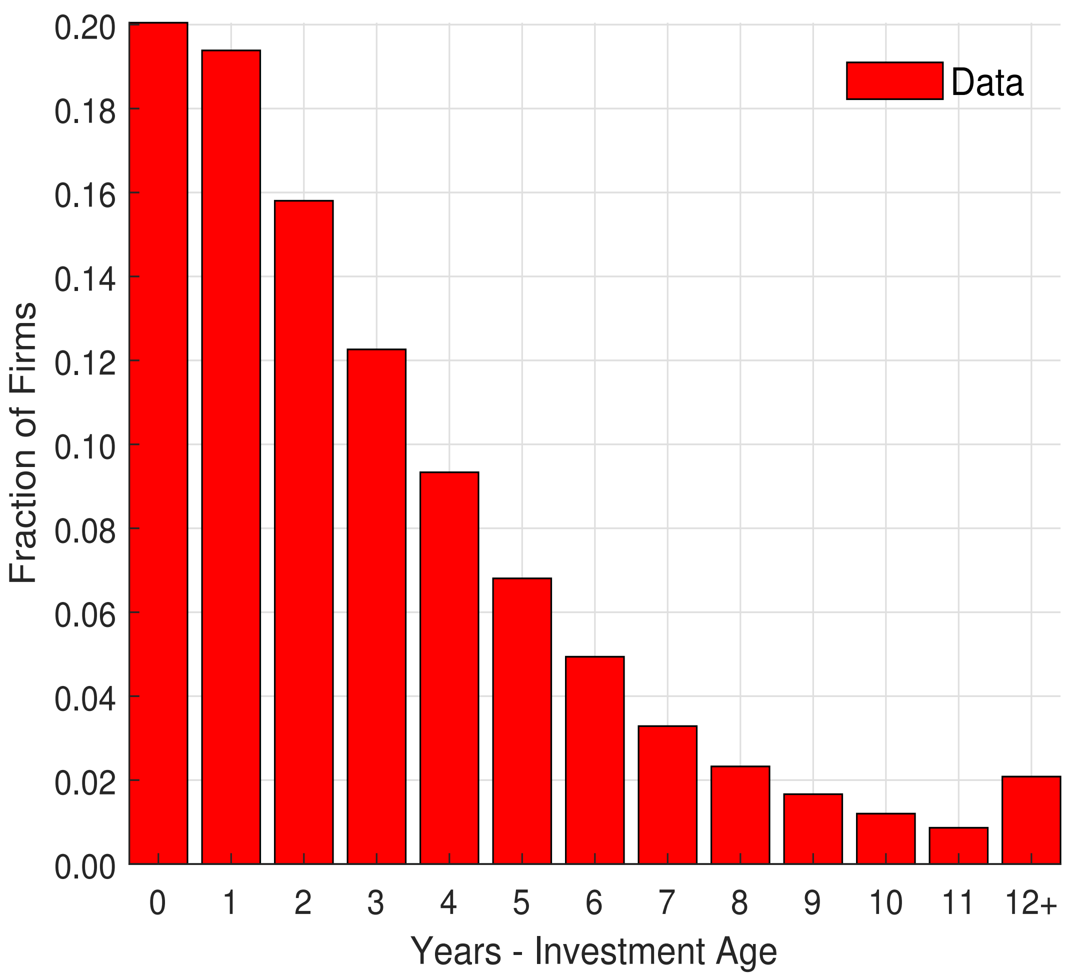

The average firm's investment cycle occurs every three and a half years. Intuitively, average investment age is negatively correlated with GDP growth. During a GDP expansion, as in the early 2000s, the share of firms that experienced an investment spike increased to about 22 per cent. During GDP contractions, as during the Global Financial Crisis in 2009 and the Sovereign Debt Crisis in 2012, more firms postponed large capital adjustments, with only about 11 percent experiencing spikes. The sharp drop in investment resulted in a significant increase in investment age and a delay in introducing new technologies in production.

Reassuringly, the investment age does not capture cohort effects related to firms' birth year, as the correlation between investment and firms' age is equal to 0.10.

3 TFP Heterogeneity and Investment Age

We now provide firm-level evidence on the negative relationship between TFP and investment age, the proxy for how advanced the technology available to the firm is—Firms that operate more obsolete technology display lower TFP. In Section 3.1, we describe the methodology employed in the analysis. In Section 3.2, we report our estimates showing that higher investment age predicts lower TFP, which point to capital-embodied technological progress and support the key mechanism at the core of the model described in Section 4.

3.1 Empirical Methodology and Specification

To estimate the empirical relationship between TFP and investment age, we fit the following specification to the panel of firms in our sample:

$$log(TFP)_{f,t} = \alpha + \beta \times Inv.Age_{f,t-1} + \Phi \times \mathbf{Controls}_{f,t} + \epsilon_{f,t}, \tag{1}$$

The dependent variable $TFP_{f,t}$ for firm f in year t is measured through the Solow residual that is computed assuming a Cobb-Douglas production function. Following Bachmann and Bayer (2014), we estimate the output elasticities of the production function as the median factor expenditures share over gross value-added within each industry. Appendix B reports additional details on the construction of the variables.

The coefficient β in equation 1 measures the sign and the magnitude of the correlation between investment age, at time *t-1*, and TFP, at time t. Consistent with the time to build in capital accumulation, we assume that it takes one period before new technologies obtained with new capital become operational in the production process.

Beyond firm's fixed effects, the set of controls features sector and sector-time dummies to account for unobserved time-variant and time-invariant industry-specific characteristics and aggregate factors that are potentially related to policy changes or business cycle fluctuations as well as sectoral trends in TFP. In sum, to estimate the empirical relationship between investment age and TFP, we exploit fluctuations in TFP around firm- and sector-specific means while simultaneously netting out common or sectoral time-varying movements of TFP across firms (through time and time-sector effects).

We interpret our estimates in a *predictive* rather than a causal sense given the identification challenges related potentially to time-varying firm-specific factors. For instance, news about current profitability or the persistence effect on firms' TFP of current shocks could induce firms to undertake a large capital adjustment. In the INVIND sample we can tackle one of the identification challenges and include in the set of control variables managers' expectations about the one-year-ahead growth rate of sales elicited in the survey at the beginning of time t, denoted by $s^e_{avg,f,t}$. Including $s^e_{avg,f,t}$ allows us to avoid confounding the role of $Inv.Age_{f,t-1}$ with otherwise unobserved future news about business prospects that may affect TFP dynamics.

3.2 Negative Relationship between Investment Age and TFP

The firm-level analysis shows that firms with higher investment age or, equivalently, less advanced technology display a lower TFP. In other words, our results support the idea that by postponing large capital adjustments, firms delay the introduction of new technologies in the production process, suffering lower TFP as the gap between the technology of investment goods currently operated by the firm and the frontier increases.

Table 2 reports estimates of the coefficient on investment age, β, in equation 1. The negative sign of the relationship and the magnitude of the coefficient are robust across data sets and control variables. Column A reports OLS results using the universe of the firms in Cerved. The negative coefficient on $Inv.Age_{f,t}$ indicates that a delay in updating capital reduces TFP, with an estimated productivity loss of about 0.9 percent per year. Column B shows that entry dynamics do not drive the results. Focusing on firms at least three years old yields a quantitatively similar coefficient. In columns C and D, we focused on the matched sample of firms in Cerved and INVIND. Column C indicates that the role of investment age for TFP dynamics is similar in magnitude when focusing on INVIND firms. Column D also includes managers' expectations about future sales, $s^e_{avg,f,t}$.[8] News about the future proxied by firms' expected sales predicts higher TFP, highlighting the role of news for productivity dynamics.

In sum, our analysis documents the predictive power of investment age, a proxy for technology, for TFP dynamics, even controlling for a large set of confounding variables, including managers' expectations about future sales. All in all, this evidence points to the role of capital embodied technological progress as a source of productivity dynamics at the firm level. Motivated by this evidence, we now formulate an equilibrium model that accounts

[8] Fiori and Scoccianti (2023) show that INVIND expectations are unbiased and informative about firms' future business prospects.

Table 2: Investment Age and Total Factor Productivity

	$TFP_{f,t}$ (A)	$TFP_{f,t}$ (B)	$TFP_{f,t}$ (C)	$TFP_{f,t}$ (D)
$Inv.Age_{f,t-1}$	-0.860^{***}	-0.837^{***}	-0.930^{***}	-0.906^{**}
	(0.00)	(0.00)	(0.01)	(0.08)
$s^e_{avg,f,t}$				0.370^{***}
				(0.02)
N. of obs.	2,756,422	2,596,223	3,347	2,773
R^2	0.746	0.750	0.919	0.918
Estimator	OLS	OLS	OLS	OLS
Firm FE	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓
Time FE	✓	✓	✓	✓
Time×Ind. FE	✓	✓	✓	✓
Source of Data	Cerved		INVIND	
Sample Period	1998-2016	1998-2016	2003-2016	2003-2016

Note: * p<0.10, ** p<0.05, and *** p<0.01, where p is the marginal probability level and is reported in parentheses. The dependent variable is the log of total factor productivity (TFP) for firm f at time t. $Inv.Age_{f,t-1}$ measures the time elapsed between investment spikes, defined as the firm experiencing an investment rate above 20 percent at time t-1. Columns A and B report estimates obtained using Cerved data. Columns C and D report estimates obtained using INVIND data that allow us in Column D to control for managers' subjective expectations about sales one year ahead.

for selected microeconomic moments about firm investment behavior and the documented firm-level evidence to study the aggregate implications of technology adoption through capital accumulation.

4 Model: Firm Heterogeneity and Technology Adoption

We formulate a general equilibrium framework featuring firm heterogeneity and investment-specific technological progress to study the aggregate implications of investment age heterogeneity. Our starting point is the neoclassical growth model of Khan and Thomas (2003, 2008), the benchmark for quantitative analysis involving firm dynamics in a general equilibrium framework. We assume capital-embodied technological progress, where the efficiency of most recent capital goods evolves stochastically around a deterministic trend as in Greenwood, Hercowitz and Krusell (2000). The firm's problem involves deciding the optimal timing for acquiring capital goods embodying the latest vintage technology. Firms that postpone acquiring the latest investment goods can invest in capital goods embodying a less efficient technology. As in Khan and Thomas (2003), the decision to purchase the latest vintage is subject to a non-convex adjustment cost contributing to reproducing the firm-level evidence on capital accumulation. Unlike our framework, in Khan and Thomas (2003), firms that postpone investing in the latest capital goods let their capital stock depreciate rather than acquiring less efficient capital goods.

In Sections 4.1 and 4.2, we outline the tradeoffs determining each firm's production and investment decision/technology adoption. Sections 4.3 and 4.4 describe the households' problem, and Section 4.5 details the recursive equilibrium of the economy. Section 4.6 discusses the model's implications for aggregate productivity.

4.1 Production

The economy consists of a continuum of firms normalized to one. As common in the literature on firm heterogeneity and investment-specific technological progress, each firm has access to an increasing and concave production function (F) that combines a predetermined stock $\tilde{k}$ of capital services rather than physical units and labor hired on a spot market to produce output y. We report model variables along the balanced growth path, deflated by their respective deterministic trends.[9] For better readability, we drop individual firms subscripts; the production function is:

$$y_t = \varepsilon_t F(\tilde{k}_t, n_t) = \varepsilon_t \tilde{k}_t^{\theta} n_t^{\nu}, \tag{2}$$

where $0 < \theta + \nu < 1$, to pin down the size of the firm. Production efficiency depends upon ε, which denotes the idiosyncratic productivity that is exogenous to the firm. As in Khan and Thomas (2008), ε is discretized so that $\varepsilon \in \{\varepsilon_1, \varepsilon_2, ..., \varepsilon_{N_\varepsilon}\}$, where $Pr(\varepsilon = \varepsilon_m | \varepsilon = \varepsilon_l) \equiv \pi^{\varepsilon}_{lm} \geq 0$, and $\sum_{m=1}^{N_\varepsilon} \pi^{\varepsilon}_{lm} = 1$ for each $l = 1, ..., N_\varepsilon$. In every period, the plant chooses its current level of employment and the stock of capital services $\tilde{k}_{t+1}$ (described in the next section) that becomes productive in the next period. Production occurs and labour is paid the real wage, denoted by w_t.

4.2 Firm's Technology Adoption and Investment Decision

The main innovation of the framework consists of modelling a vintage technology structure for capital-embodied technological change. The stock of capital *services* evolves according to the following

$$\gamma_k \tilde{k}_{t+1} = (1-\delta)\tilde{k}_t + \tilde{i}_{j,t}, \tag{3}$$

where $\delta \in (0, 1)$ is the rate of physical capital depreciation and γ_k is the deterministic gross growth rate of capital services along the balanced growth path. As in Solow (1960) and Greenwood, Hercowitz and Krusell (1997, 2000), we formalize the notion of investment-specific technological change by assuming that the inflow of new capital services, $\tilde{i}_{j,t}$, is the product of the physical quantity of investment good, $i_{j,t}$, and its associated vintage produc-

[9] Along the balanced growth path, γ_q denotes the gross trend growth rate of the investment-specific technology frontier. Output, consumption, and investment grow at a gross rate of $\gamma = \gamma_q^{\theta/(1-\theta)}$, while capital grows at a gross rate of $\gamma_k = \gamma_q^{1/(1-\theta)}$.

tivity, denoted with a slight abuse of notation by $q_{j,t}$.[10] Alternatively, $1/q_{j,t}$ can be interpreted as representing the cost of producing one unit of capital (in terms of final output) of a given vintage j. The index j takes nonnegative values and identifies the current vintage of technology available to each firm and is optimally chosen by the firm, with $j = 0$ denoting the latest vintage and higher values of j indicating older vintages.

Note that $\tilde{k}$ denotes the stock of *capital services* obtained with investment goods and therefore does not have a subscript j. As in Khan and Thomas (2003) and in the spirit of Solow (1960) and the subsequent literature, expressing the stock of capital services in units of efficiency rather than physical units allows for a convenient aggregation of physical units of investment embodying different levels of efficiencies, accumulated by the firm over time. The investment-specific technological frontier grows deterministically at the gross rate of $\gamma_q > 1$. Along the balanced growth path, the efficiency of the latest vintage $q_{0,t}$ fluctuates around its deterministic trend and follows a Markov process in logs. Similarly to ε, $q_{0,t}$ is discretized so that in every period $q_0 \in \{q_{0,1}, q_{0,2}, ..., q_{0,N_q}\}$, where $Pr(q_0 = q_{0,z}|q_0 = q_{0,n}) \equiv \pi_{nz}^{q_0} \geq 0$, and $\sum_{z=1}^{N_{q_0}} \pi_{nz}^{q_0} = 1$ for each $n = 1, ..., N_{q_0}$.

The timing of events is as follows. Each firm starts the period with a predetermined stock of capital services, $\tilde{k}_t$, and a given vintage, $q_{j,t}$. The idiosyncratic productivity ε_t and the efficiency of the capital-embodied technology $q_{0,t}$ are realized, and firms draw the identically and independently distributed cost of adopting the latest vintage ξ. Then, firms decide whether to pay the cost to switch technology by acquiring the latest investment good and choosing the capital stock for the next period. Each unit of capital acquired comes at a cost of $1/q_{0,t}$. A firm that pays the adjustment cost adopts the latest vintage in the current period so that $\tilde{k}_{t+1}$ depends on the efficiency $q_{0,t}$ and the quantity of investment $i_{0,t}$ undertaken. A firm that postpones adopting the latest vintage chooses $\tilde{k}_{t+1}$. In this case, the firm keeps

[10]Note that later we refer to $q_{j,t}$ as the vintage at the beginning of the period and not to the one chosen by the firm at time t.

acquiring investment goods embodying a technology that becomes less efficient relative to the deterministic trend at a rate γ_q $(q_{j,t}/\gamma_q)$. At the beginning of the next period, the efficiency of the adopters of the latest vintage is $e^{\rho_q \times log(q_{0,t})}$ as $q_{0,t}$ reverts to its mean at a rate of ρ_q. Similarly, firms that have postponed adjusting start the next period with a vintage that has efficiency equal to $q_{j,t}/\gamma_q$. To be clear, firms not adopting the latest vintage keep acquiring investment goods embodying a constant *level* of efficiency, but, along the balanced growth path and relative to the frontier, become less efficient given the availability of newer investment goods embodying a more efficient technology. In the next section, we describe the tradeoffs associated with the choice of technology adoption.

As in Khan and Ravikumar (2002), the adoption adjustment cost ξ is non-convex, and its modeling strategy follows Caballero and Engel (1999) and the subsequent literature on lumpy investment. Thus, the decision to adopt the latest vintage involves a non-convexity; conditional on adjusting capital and upgrading technology, the cost ξ incurred is independent of the scale of adjustment. To be sure, the investment size in the latest vintage determines the increase in the stock of capital services. As in Thomas (2002) and the subsequent literature, we assume that ξ is independently and identically distributed (i.i.d.) across firms and time.

In each period, a firm is defined by its vintage productivity q_j, its idiosyncratic productivity level $\varepsilon \in \mathcal{E} \equiv \{\varepsilon_1, \varepsilon_2, ..., \varepsilon_{N_\varepsilon}\}$, its predetermined stock of capital $\tilde{k}_t \in \mathbf{R}_+$, and its cost associated with vintage adoption $\xi \in [\overline{\xi}_l, \overline{\xi}_h]$, which is denominated in units of labor and drawn from the time-invariant distribution G common to all production units and assumed to be uniform. As the firm's current adjustment cost has no implications for its future adjustment, the distribution of firms is summarized by $(\varepsilon_t, q_{j,t}, \tilde{k}_t)$. To characterize the distribution of firms over $(\varepsilon_t, q_{j,t}, \tilde{k}_t)$, we use the probability measure μ defined on the Borel algebra S for the product space $S = \mathcal{E} \times \mathbf{R}_+ \times \mathbf{R}_+$. The aggregate state of the economy is described by

$(q_{0,t}, \mu_t)$, the efficiency of the technology frontier and the distribution of firms that evolves according to a mapping $\Gamma : \mu_{t+1} = \Gamma(q_{0,t}, \mu_t)$.

We note that when the cost of technology adoption is equal to zero in every period, the model becomes a standard neoclassical growth model with firm heterogeneity and investment-specific technological progress.

4.3 Firm's Dynamic Programming Problem

To describe the adoption and the investment decision of the firm, as is customary in the literature, we adopt the approach of Khan and Thomas (2008) and state the problem in terms of utils of the representative households (rather than physical units), and denote the marginal utility of consumption by $p_t = p(q_{0,t}, \mu_t)$. This variable indicates the pricing kernel used by firms to price output streams. Given the i.i.d. nature of the adjustment cost ξ_t, continuation values can be integrated out of future continuation values.

Let $V^1(\varepsilon_t, q_{j,t}, \tilde{k}_t, \xi_t; q_{0,t}, \mu_t)$ denote the expected discounted value of a firm entering the period with $(\varepsilon_t, q_{j,t}, \tilde{k}_t)$ and drawing an adjustment cost ξ when the aggregate state of the economy is $(q_{0,t},\ \mu_t)$. The dynamic optimization problem for the typical firm is described using a functional equation defined by equations 4, 5, and 6. First, we define the beginning-of-period expected value of a firm before the realization of its fixed cost draw but after the determination of $(\varepsilon_t, q_{j,t}, \tilde{k}_t)$:

$$V^0(\varepsilon_t, q_{j,t}, \tilde{k}_t; q_{0,t}, \mu_t) = \int_{\overline{\xi_l}}^{\overline{\xi_h}} V^1(\varepsilon_t, q_{j,t}, \tilde{k}_t, \xi_t; q_{0,t}, \mu_t) dG(\xi). \tag{4}$$

The firm's profit-maximization problem, which takes as given the evolution of the firm distribution, $\mu_{t+1} = \Gamma(q_{0,t}, \mu_t)$, is then described by

$$V^1(\varepsilon_t, q_{j,t}, \tilde{k}_t, \xi_t; q_{0,t}, \mu_t) = \max_{n_t, k^A_{t+1}, k^{NA}_{t+1}} \left\{ \begin{array}{l} \left[\varepsilon_t F(\tilde{k}_t, n_t) - w(q_{0,t}, \mu_t) n_t\right] p(q_{0,t}, \mu_t) + \\ \max \left[\begin{array}{l} \frac{(1-\delta)\tilde{k}_t}{q_{0,t}} p(q_{0,t}, \mu_t) - \xi_t w(q_{0,t}, \mu_t) \\ +R(\varepsilon_t, q_{0,t}, \tilde{k}^A_{t+1}; q_{0,t}, \mu_t); \\ \frac{(1-\delta)\tilde{k}_t}{q_{j,t}/\gamma_q q_{0,t}} p(q_{0,t}, \mu_t) + R(\varepsilon_t, q_{j,t}, \tilde{k}^{NA}_{t+1}; q_{0,t}, \mu_t) \end{array} \right] \end{array} \right\} \tag{5}$$

s.t. n_t, $\tilde{k}^A_{t+1}$, and $\tilde{k}^{NA}_{t+1} \in \mathbf{R}_+$,

where $R(\varepsilon_t, q_{j,t}, \tilde{k}^A_{t+1}; q_{0,t}, \mu_t)$ represents the continuation value associated with a given combination of the idiosyncratic shock, the vintage $q_{0,t}$, and the stock of capital in efficiency units $\tilde{k}_{t+1}$. For firms willing to pay the fixed cost, $R(.)$ is the following:

$$\begin{aligned} R(\varepsilon_l, q_{j,t}, \tilde{k}_{t+1}; q_{0,t}, \mu_t) \equiv -\frac{\gamma_k \tilde{k}_{t+1}}{q_{0,t}} p(q_{0,t}, \mu_t) \\ +\beta \sum_{m=1}^{N_\varepsilon} \pi^{q_0}_{nz} \pi^{\varepsilon}_{lm} V^0(\varepsilon_m, e^{\rho_q log(q_{0,t})}, k_{t+1}; q_{0,t+1}, \mu_{t+1}), \end{aligned} \tag{6}$$

where firms buy capital goods that embody the latest vintage $q_{0,t}$ at time t and have a continuation value at the beginning of the period evaluated over the possible realizations of the idiosyncratic productivity ε, the vintage considering that $q_{0,t}$ reverts to its mean at a rate ρ_q and the capital $\tilde{k}_{t+1}$, as well as the aggregate states. For firms that postpone paying the adjustment cost ξw, the current vintage is $q_{j,t}/\gamma_q$ that is carried forward into the next period. For notational convenience, as in Khan and Thomas (2008), rather than subtracting investment from current profits, the value of undepreciated capital augments current profits, and the firm is seen to repurchase its capital stock each period as this approach is equivalent but notationally more convenient. Thus, we let $Q(\varepsilon, q_j, \tilde{k}, \xi; q_0, \mu)$ and $K(\varepsilon, q_j, \tilde{k}, \xi; q_0, \mu)$ represent the choice of technology and capital for the next period by firms of type $(\varepsilon, q_j, \tilde{k})$ with adjustment cost ξ, while $N(\varepsilon, q_j, \tilde{k}, \xi; q_0, \mu)$ denotes firms' labor choices.

4.4 Households

The economy features a continuum of identical households that consume the output good and supply labor to all firms. Households maximize their lifetime expected utility W by choosing current consumption c and labor effort N^h:

$$W(\lambda_t; q_{0,t}, \mu_t) = \max_{c_t, N_t^h, \lambda_{t+1}} \left[U\left(c_t, N_t^h\right) + \beta \sum_{z=1}^{N_{q_0}} \pi_{nz}^{q_0} W(\lambda_{t+1}; q_{0,t+1}, \mu_{t+1}) \right] \tag{7}$$

Households can trade one-period shares in all firms, denoted by the measure λ_t, and have access to a complete set of state-contingent claims. As there is no household heterogeneity assets are in zero net supply in equilibrium. Define

$$p(q_{0,t}, \mu_t) \equiv U_C\left(c_t, N_t^h\right). \tag{8}$$

The labor supply schedule is defined by the first-order condition that equates the marginal rate of substitution between hours and consumption to the real wage w:

$$p(q_{0,t}, \mu_t) w(q_{0,t}, \mu_t) = U_{N_t^h}\left(c_t, N_t^h\right), \tag{9}$$

Let $C(\lambda_t; q_{0,t}, \mu_t)$ describe the household choice of current consumption, $N^h(\lambda_t; q_{0,t}, \mu_t)$ the current allocation of time working, and $\Lambda^h(\varepsilon_{t+1}, q_{j,t+1}, \lambda_t, \tilde{k}_{t+1}; q_{0,t}, \mu_t)$ the quantity of shares purchased in plants that begin the next period with productivity ε_{t+1}, vintage $q_{j,t+1}$, and $\tilde{k}_{t+1}$ capital services.

4.5 Recursive Equilibrium

A recursive competitive equilibrium is a set of functions $(p,\ v^1,\ Q,\ K,\ W,\ C,\ N,\ N^h,\ \Lambda^h,\ \Gamma)$ that satisfy the firms' and households' problem and clear the markets for assets, labor, and output:

(i) Firm's optimality: Taking $p_t = p(q_{0,t}, \mu_t)$ as given, $V^1\left(\varepsilon_t, q_{j,t}, \tilde{k}_t, \xi_t; p_t\right)$ solves equations 4, 5, and 6 and the corresponding policy functions $N = N\left(\varepsilon_t, q_{j,t}, \tilde{k}_t, \xi_t; p_t\right)$, $Q = Q\left(\varepsilon_t, q_{j,t}, \tilde{k}_t, \xi_t; p_t\right)$, and $K = K\left(\varepsilon_t, q_{j,t}, \tilde{k}_t, \xi_t; p_t\right)$.

(ii) Household's optimality: Taking p_t as given, the household's decisions satisfy equations 7 and 9 and the corresponding policy functions (C, N^h, Λ^h).

(iii) $\Lambda^h(\varepsilon_{l,t+1}, q_{j,t+1}, \tilde{k}_{t+1}, \mu_t; q_{0,t}, \mu_t) = \mu(\varepsilon_{l,t+1}, q_{j,t+1}, \tilde{k}_{t+1})$ for each $(\varepsilon_{l,t+1}, q_{j,t+1}, \tilde{k}_{t+1}) \in S$.

(iv) Commodity market clearing: $C_t = \int y_t d\mu - \int \int_{\bar{\xi}_l}^{\bar{\xi}_h} [[\gamma_k K\left(\varepsilon_t, q_{j,t}, \tilde{k}_t, \xi_t; q_{0,t}, p\right) - (1-\delta) K] / Q\left(\varepsilon_t, q_{j,t}, \tilde{k}_t, \xi_t; q_{0,t}, p\right)] dG d\mu$.

(v) Labor market clearing: $N^h = \int N\left(\varepsilon_t, q_{j,t}, \tilde{k}_t; q_{0,t}, \mu_t\right) d\mu + \int_{\bar{\xi}_l}^{\bar{\xi}_h} \xi J(x) dG d\mu$, where $J(x) = 0$, if the firm does not upgrade its vintage, and 1 otherwise.

(vi) Model-consistent dynamics: The evolution of the cross-sectional distribution that characterizes the economy, $\mu_{t+1} = \Gamma(q_{0,t}, \mu_t)$, is induced by the adjustment decision and the exogenous processes for ε. Conditions (i), (ii), (iii), (iv), and (v) define an equilibrium given Γ, while condition (vi) determines the equilibrium condition for Γ. We confine to Appendix C the discussion about the (S,s) decision rule for the firm upgrading and investing decision and the details on the evolution of the cross-sectional distribution of firms' productivity and capital stocks.

4.6 Cross-Sectional Distribution of Technology and Investment

In this section, we discuss the role of investment-specific productivity heterogeneity for aggregate dynamics. Non-convex adoption costs imply that the firm's technology adoption/investment decision follows an (S,s) rule: Some firms invest in capital goods embodying the latest vintage technology, while others postpone it and keep investing in investment goods embodying the current vintage they operate.[11] The efficiency gap between the vintage currently available to the firm and the technological frontier, together with the realization of the idiosyncratic and aggregate shocks affect the timing of technology adoption at the firm level. Thus, in equilibrium, firms invest in vintages of different qualities. The history of adoption choices contributes to investment-specific heterogeneity determining the *economy-wide* productivity. Shifts in the cross-sectional distribution, which are determined by variations in the firms' adoption decision and are driven by shocks, result in aggregate investment-specific productivity fluctuations that constitute an additional force in the propagation of aggregate shocks shaping investment and therefore output dynamics. The main feature distinguishing our other frameworks with investment-specific technological change from those in Solow (1960), Greenwood, Hercowitz and Krusell (1997, 2000) and Khan and Thomas (2003), beyond the presence of idiosyncratic productivity shocks, is that firms can decide whether to invest in the latest vintage of investment goods or to keep buying investment goods embodying a less efficient technology relative to the frontier. However, in Solow (1960), Greenwood, Hercowitz and Krusell (1997); **?**, 2000), and Khan and Thomas (2003) only the latest vintage is available. In the presence of heterogeneity in investment-specific productivity, firms' investment responds not only to aggregate shocks that alter the timing of adoption but also to productivity gains as firms move from obsolete technology to the latest and most efficient technology.

[11] See Appendix C for additional details.

Before describing the calibration strategy of the model, it is worth examining the role of three key parameters. The support of the adjustment cost distribution ($\overline{\xi}_l, \overline{\xi}_h$) determines the magnitude of the cost of adjusting technology. The higher value of the adjustment cost, the higher the potential cost of adopting the latest vintage. Increasing this cost leads to a higher average investment age increasing the mass of firms experiencing spikes. The idiosyncratic process's persistence and standard deviation interact with the vintage effect in shaping the economic incentives that make the firm adopt the latest vintage and choose the appropriate level of capital.

5 Taking the Model to the Data

We now take the model to the data. In Section 5.1, we describe the parameterization of the model, and in Section 5.2 we describe the solution algorithm.

5.1 Parameterization

Following the business cycle literature, we select parameters to fit critical micro and macro first-order moments of the Italian economy. Table 3 summarizes parameter values, targeted moments, and data sources. We are to assign values to 14 parameters related to the growth rate of investment-specific technological progress and aggregate variables along the balanced growth path in the absence of shocks (γ_q, γ, and γ_k), the production process (δ, θ, and ν), individual preferences on disutility of labor and the discount factor (A and β), the technology adjustment cost function ($\overline{\xi}_l$,$\overline{\xi}_h$), the exogenous idiosyncratic productivity process (ρ_ε, and σ_ε), and the evolution of the technological frontier around its deterministic trend (ρ_q, and σ_q). We first describe the set of parameters that are calibrated using independent evidence. Then, we focus on those estimated within the model to reproduce relevant targets in the data.

Parameter values based on a priori information. One period in the model represents one year, corresponding to the data frequency employed in Sections 2 and 3. The depreciation rate is taken from the Italian National Institute of Statistics and is equal to 9 percent. The elasticity of output to capital (θ) and labor (ν) in the production process is set rqual to 0.18 and 0.64, respectively, the median of the sector-level estimates employed to construct TFP in Section 3.1. To construct a measure of investment-specific technological progress, we follow Khan and Thomas (2003) in adapting the procedure in Christiano and Fisher (1998). Based on the price series for aggregate investment relative to the one for nondurables and services consumption expenditures, we can use Italian data to directly estimate γ_q, ρ_q, and σ_q. Given the gross growth rate of the technological frontier γ_q and θ, balanced growth rate restrictions pin down the growth rate of output, consumption, and investment, $\gamma = \gamma_q^{\theta/(1-\theta)}$, as well as the one for capital $\gamma_k = \gamma_q^{1/(1-\theta)}$.

Parameter values based on average Italian data. We follow the literature on firm heterogeneity and assume a perfectly elastic labor supply as in Hansen (1985) and Thomas (2002), with the disutility of labor, A, set equal to 1.23 to reproduce employment equal to 0.6. The discount factor β is set equal to 0.975 to reproduce the real annual interest rate in the data.

The remaining parameters are calibrated to match targeted moments in the data. While none of the parameters have a one-to-one relationship to a specific moment, it is instructive to describe the calibration as a few distinct steps. The lower and the upper support of the technology adjustment cost function ($\overline{\xi}_l$,$\overline{\xi}_h$) are set to reproduce the share of positive investment accounted for by firms experiencing spikes and the average investment age of the firms' cross-sectional distribution, respectively. The estimated $\overline{\xi}_l$ and $\overline{\xi}_h$ imply that adjustment costs, conditional on adjustment, are equal to about 3 percent of annual output, which is on the lower end of the previous estimates in the existing literature; see, for instance,

Table 3: Benchmark Calibration

Parameter		Value	Target
Depreciation rate	δ	0.091	Data
Elasticity of output w.r.t. capital	θ	0.18	Data
Elasticity of output w.r.t. labor	ν	0.64	Data
Mean growth rate of technology frontier	γ_q	1.0081	Data
Mean growth rate of output	γ	1.002	Balanced growth path restriction
Mean growth rate of capital	γ_k	1.010	Balanced growth path restriction
Persistence latest technology	ρ_q	0.55	Data
St. dev. latest technology	σ_q	0.006	Data
Disutility of labor	A	1.234	Employment rate = 60%
Discount factor	β	0.975	Annual real interest rate = 2.3%
Persistence exogenous idios. productivity	ρ_ε	0.801	TFP persistence in the data
St. dev. idiosyncratic productivity	σ_ε	0.018	Sh. of firms experiencing $ik > 0.20$
Lower bound adoption cost function	$\bar{\xi}_l$	0.145	Sh. of pos. invest. by firms with $ik > 0.20$
Upper bound adoption cost function	$\bar{\xi}_h$	0.651	Average $Inv.Age$

Bachmann, Caballero and Engel (2013).[12] The persistence of the exogenous idiosyncratic productivity process (ρ_ε) is set so that the model firms' TFP displays the same persistence as the one estimated in the data (0.80). This coefficient obtains fitting an autoregressive process of order one to the logarithm of firm-level TFP.[13] The standard deviation of the idiosyncratic productivity process (σ_ε) is selected to reproduce the fraction of firms experiencing an investment spike in the data. In Section 7, we compare the business cycle performance of the vintage model with a standard neoclassical benchmark with firm heterogeneity. Our vintage model nests the neoclassical benchmark as a special case that is obtained when the technology adjustment cost is equal to zero and there is no technological regress.

[12]The magnitude of the estimated costs in our framework is on the lower end of previous estimates as shown in Table 4 on page 47 in Bachmann, Caballero and Engel (2013). Our estimates are similar in magnitude to Bachmann, Caballero and Engel (2013)—adjustment costs equal to 3.6 percent—and Khan and Thomas (2008),between 0.5 and 3.7 percent. Bloom (2009) and Cooper and Haltiwanger (2006) report adjustment costs of about 30 percent of firms' output.

[13]Adding firm-specific fixed effects to the regression drops the autoregressive coefficient to 0.45. Adopting a lower persistence of the idiosyncratic process does not affect quantitatively the results of the paper.

5.2 Solution Algorithm

We follow the approach in Khan and Thomas (2008) to solve the model. This strategy replaces the aggregate law of motion for the distribution with a forecast rule. Typically, to predict prices and the future proxy aggregate state, agents use the mean capital stock. Our framework features two endogenous cross-sectional distributions for the capital stocks and the vintage technologies.[14] As a result, we proxy the distribution with the mean capital stock, $\overline{K}_{t+1}$, and the mean of the investment-specific productivity distribution, $\overline{Q}_{t+1}$, and their interaction. We use the same set of regressors to forecast the marginal utility of consumption, p_t. Moreover, we estimate the rule conditional on each realization of the aggregate process, $q_{0,t}$. We report additional details regarding the solution method in Appendix D.1.

6 Microeconomic Heterogeneity in the Model and the Data

We now discuss the ability of the model to fit multiple dimensions of the microeconomic heterogeneity in the data related to firms' investment behavior and TFP. In Sections 6.1 and 6.2, we show that the model fits well the cross-sectional distribution of investment rates and the timing of investment spikes across firms—i.e., the empirical proxy for vintage technology. In Section 6.3 , we discuss the relationship between dispersion in idiosyncratic investment-specific productivity in the data and in the model.

6.1 Cross-Sectional Distribution of Investment Rates

We now examine the model's performance, starting from the cross-sectional distribution of investment rates. As in Cooper and Haltiwanger (2006) and Khan and Thomas (2008), the distribution is summarized using five groups: inaction, positive and negative investment,

[14] We note that, even when two firms have the same TFP, they may still choose a different stock of capital next period, as the continuation value depends upon ε_t and $q_{j,t}$.

and positive and negative spikes. A specific threshold for the investment rate (*ik*) identifies each group. Before examining the results in Table 4, we note that our definition of the inaction region is broader than the definition employed in the existing theoretical literature and more in line with the empirical literature—see Øivind and Schiantarelli 2003. This choice allows us to capture the small investment rates occurring in about one-third of the firms in the sample.

Table 4: Distribution of Firm Investment Rates

	Inaction $\lvert ik \rvert \leq 0.05$ (A)	Positive Spikes $ik > 0.20$ (B)	Negative Spikes $ik < -0.20$ (C)	Positive Investment $ik > 0.05$ (D)	Negative Investment $ik < -0.05$ (E)
Data	37.21%	20.04%	2.32%	57.99%	4.80%
Baseline Vintage	24.30%	20.04%	0.61%	61.81%	14.89%
Neoclassical $\bar{\xi}_{l,h} = 0$	20.68%	22.26%	0.00%	78.31%	1.01%

Note: Each entry reports the fraction of firms that every year, on average, exhibit investment rates that fall in each category. The "Neoclassical $\bar{\xi}_{l,h} = 0$" model retains the parameters of the baseline vintage model except for $\bar{\xi}_l, \bar{\xi}_h = 0$.

The average share of firms experiencing a positive spike is matched by virtue of the calibration strategy discussed in Section 5.1. The model closely matches the fraction of firms experiencing negative spikes and positive investment. However, the model overstates the share of firms downsizing capital and understates the share of inactive firms. A comparison with a neoclassical benchmark sheds light on the role of vintage technology in driving aggregate dynamics. Specifically, we consider the model "Neoclassical $\bar{\xi}_{l,h} = 0$" that retains the same set of parameters as the vintage model and sets the adjustment cost to technology, $\bar{\xi}_l$ and $\bar{\xi}_h$ equal to zero.

Relative to the neoclassical benchmark ($\bar{\xi}_{l,h} = 0$), vintage effects shift the cross-sectional distribution of investment rates by reducing the capital chosen by firms. The reason behind this is that the realizations of the idiosyncratic process and the distance from the vintage frontier have conflicting effects on the adoption decision of the firm. Unfavorable realizations of the idiosyncratic process tend to make firms postpone the adoption decision. However, delaying the adoption of the latest vintage increases the distance from the productivity frontier. Therefore, firms reduce their capital stock more than they would have, absent the vintage effect.

The vintage model also reproduces the persistence of firm-level investment rates in the data. As known in the literature since Caballero and Engel (1999), firm-level investment rates exhibit nearly zero autocorrelation. Our model successfully replicates this feature of the data with a 0.03 autocorrelation—close to 0.10 in the data. In the neoclassical models, firm-level investment rates are negatively correlated, failing to match this feature of the data with autocorrelation equal to -0.10 ($\overline{\xi}_{l,h}$=0).

6.2 Cross-Sectional Distribution of Investment Age

We now describe the implied cross-sectional distribution of investment age in the model and the data. As in the data, investment age is calculated as the time elapsed since the firm experienced a positive investment spike. The cross-sectional distribution of investment age displays a sizable and negative correlation (-0.77) with the cross-sectional distribution of investment-specific productivity and, thus, conveys information about vintage technology. The opposite is true for the neoclassical model that displays a 0.02 correlation.

We remind the reader that our calibration targets the average investment age of about 3 and a half years and, as discussed in the previous subsection, the fraction of firms experiencing an investment spike—i.e., firms exhibiting an investment age equal to zero. As shown in Figure 2 the model distribution closely matches the data, with a similar decay in invest-

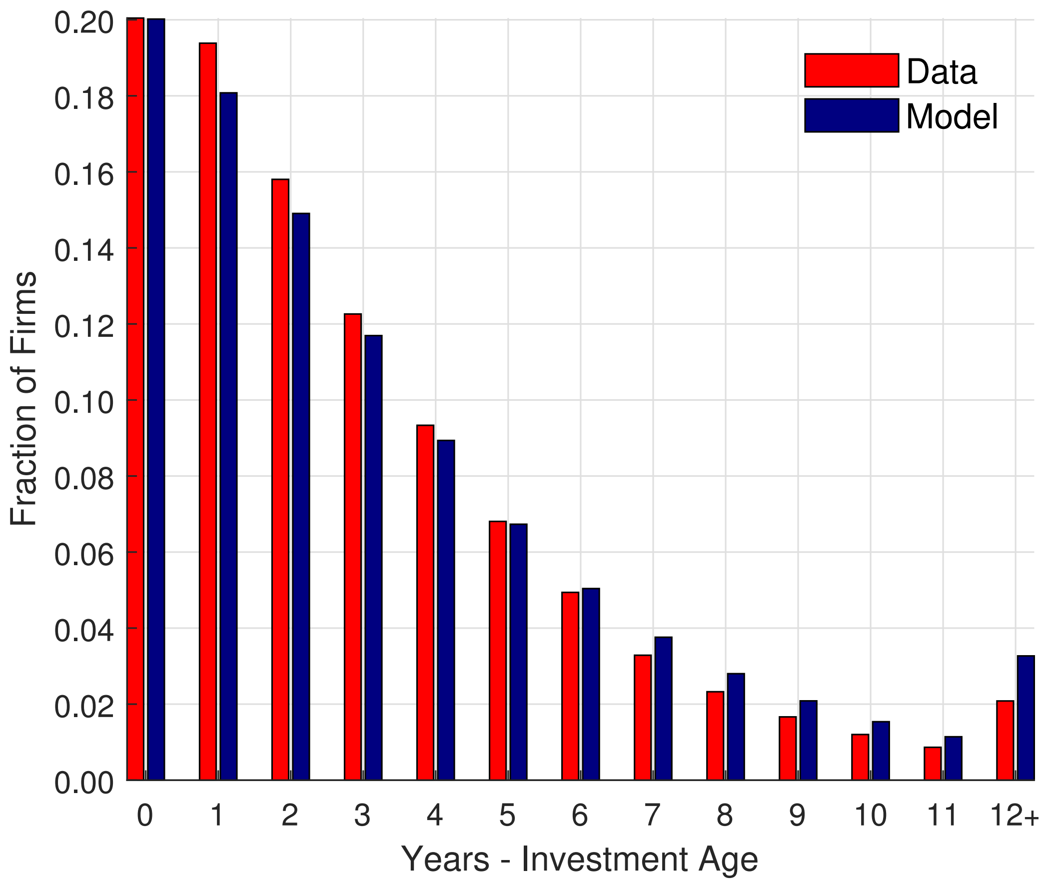

Figure 2: Comparison between the Empirical and the Model-Based Investment Age Distribution

ment age. Overall, the performance of the model is satisfactory. We notice that the ability of the framework to reproduce the timing of investment spikes across firms is distinct from the model's success in accounting for the cross-section of investment rates. While the fraction of firms with an investment age of zero coincides by construction with the fraction of firms exhibiting spikes, the fraction of firms with an age of one and above depends on the investment behavior of firms—determined by the realization of individual states.

For comparison, setting the upper support of the nonconvex adjustment cost of technology adoption equal to zero, reduces the average investment age of the model to 2.95 years.

6.3 Model and Data: The Cross-Sectional Distribution of TFP

We now discuss how heterogeneity in investment-specific productivity contributes to dispersion in *TFP* across firms. As discussed in Greenwood, Hercowitz and Krusell (1997), the neoclassical growth model with investment-specific technical progress can be transformed in a model with neutral technological change. In this case, the capital stock is measured in terms of consumption (rather than units of efficiency), and the relative price is constant and equal to one. Under this transformation, the measured Solow residual combines neutral and investment-specific productivity. Specifically, applying the same transformation in Greenwood, Hercowitz and Krusell (1997), we show that the total factor productivity for each firm f can be expressed as $TFP_{f,t} = \varepsilon_t q^{\theta}_{j,t-1}$.[15] Using this formula, we compute TFP for each firm, and the resulting cross-sectional dispersion in the vintage model and in the neoclassical benchmark with adoption cost $\bar{\xi}$ equal to zero.

To assess the contribution of vintage technology to the cross-sectional dispersion, we compute the log of the 90/10 range of firm-level TFP in the model and compare it with Cerved data. The latter obtains as dispersion of residuals from regressing the log of TFP on sector, time, and sector-time effects over the 1998-2016 sample, the same specification used in the empirical analysis in Section 3.

Two main results emerge from Table 5 that highlight the quantitative importance of investment-specific vintage technology as a source of TFP heterogeneity in the data. First, the vintage technology model accounts for about 10 percent of the TFP dispersion in the data. We interpret the estimated contribution of the vintage structure to TFP heterogeneity as a lower bound given that there are no spillovers from the investment-specific productivity to the "pure" neutral component of TFP, with ε assumed to be exogenous. Second,

[15] Under the mentioned transformation, the law of motion of capital features economic rather than physical depreciation: $k_{t+1} = (1-\tilde{\delta})k_t + i_t$, where $(1-\tilde{\delta}) = (1-\delta)\frac{q_{t-1}}{q_t}$; see, for a discussion, Appendix B on page 360 in Greenwood, Hercowitz and Krusell (1997).

Table 5: Data and Model: TFP Dispersion

	Data	Baseline Vintage	Neoclassical $\bar{\xi}_{l,h} = 0$
	(A)	(B)	(C)
$\log(TFP_{90}/TFP_{10})$	1.00	0.09	0.06
<u>Decomposition</u>			
$\log(q_{j,90}/q_{j,10})$	n.a.	0.09	0
$\log(\varepsilon_{90}/\varepsilon_{10})$	n.a.	0.06	0.06

Note: The $\log(TFP_{90}/TFP_{10})$ is the log of the ratio between the 90^{th} and 10^{th} percentiles of the cross-sectional distribution of TFP. Similarly, $\log(q_{j,90}/q_{j,10})$ and $\log(\varepsilon_{90}/\varepsilon_{10})$ denote the 90/10 range of investment-specific productivity and exogenous idiosyncratic productivity, respectively. The "Neoclassical $\bar{\xi}_{l,h} = 0$" model retains the parameters of the baseline vintage model except for $\bar{\xi}_l, \bar{\xi}_h = 0$.

heterogeneity in investment-specific productivity q_j amplifies TFP dispersion relative to the neoclassical model by about 50 percent. Note that the contribution to TFP heterogeneity in dispersion in the log of q_j is multiplied by θ, the exponent of capital services in the production function.

Moreover, the vintage model generates a negative correlation between TFP dispersion and output (-0.58). When the efficiency of the latest vintage grows faster than expected, TFP dispersion declines because more firms adopt new technology to avoid operating a more obsolete technology. In the literature, the cyclicality of the cross-sectional dispersion of TFP proxies aggregate uncertainty shocks; see, for instance, Bloom et al. (2018). In our framework, the result is obtained because of the endogenous timing of investing in new technologies.

7 Vintage Technology Amplifies Aggregate Fluctuations

In this section, we show that the microeconomic heterogeneity in the efficiency of investment goods amplifies the magnitude of business cycles statistics in response to technology shocks relative to a neoclassical benchmark. In Section 7.1 we report business cycle statistics, and in Section 7.2 we report the impulse response functions.

7.1 Business Cycle Properties

To assess the role of vintage technology for the business cycle dynamics, we compare the properties of the aggregate series obtained by simulating the baseline vintage model with a neoclassical benchmark—i.e., the vintage model with zero adoption cost $\overline{\xi_l}, \overline{\xi_h} = 0$. Unlike the neoclassical growth model, in which firms can always invest in goods embodying the latest vintage technological frontier, idiosyncratic and aggregate technology shocks alter the timing of technology adoption and affect each firm's production possibility frontier. As a result, the distribution of technology across firms determines the aggregate efficiency of the economy in converting investment into capital services.

In the model, the investment-specific technology frontier evolves stochastically around a deterministic trend, technology innovations directly affect only the adopters of the latest vintage, increasing the distance from the frontier for nonadopters. We perform the conventional business cycle exercise by simulating the model following investment-specific technology shocks to assess the amplification in aggregate dynamics due to investment-specific technological progress.[16]

Table 6 reports the canonical HP-filtered business cycle statistics of the aggregate series

[16] We discretize the aggregate technological process so that the realizations of the shock are such that there is no technological regress—i.e., the growth rate of technological efficiency, $q_{0,t}$, is always non-negative. Removing this assumption does not significantly affect the main results of the paper.

in the data and the counterparts for the baseline vintage and the neoclassical models.

The key message is that the vintage model improves the propagation mechanism of investment-specific technology shock and amplifies the volatility of the aggregate series relative to the neoclassical benchmarks considered. In response to the same shocks, the standard deviation of GDP in the vintage model nearly doubles relative to the neoclassical benchmark.

The amplification of aggregate dynamics depends on the joint, reinforcing effect of investment and TFP.[17] First, aggregate shocks alter the incentives to adopt new technologies, shifting the cross-sectional distribution of investment-specific productivity and contributing to aggregate fluctuations in the efficiency of investment and TFP, over and above the initial impulse of the stochastic technological frontier. The larger response of more efficient investment and employment, in turn, feeds to output. Unlike the standard model with investment-specific shocks, the technology adoption framework can generate a positive comovement between GDP and consumption. Specifically, the output response is strong enough to lift consumption and induce a positive correlation between the two series. Moreover, the vintage model delivers a persistence of the cyclical series that is higher than the standard neoclassical model, indicating that technology adoption significantly improves the endogenous propagation of technology shocks. It is now well understood in the literature that a single shock, plausibly calibrated, cannot account for the entire business cycle dynamics. In the vintage model, calibrated investment-specific shocks account for around a third of the GDP volatility in the data and about 60 percent of the fraction of the cyclical investment fluctuations, a result in line with Smets and Wouters (2007) and Justiniano, Primiceri and Tambalotti (2010).[18]

[17] We define aggregate TFP as the weighted average of the firm-specific $TFP_{f,t}$=$\varepsilon_{f,t} \times q^{\theta}_{j,t}$.

[18] The rest of GDP volatility is likely accounted by other shocks to policy (monetary or fiscal), uncertainty, or competition, just to name a few.

Table 6: Business Cycle Statistics: Technology Shock

	GDP	Consumption	Investment	Employment
	(A)	(B)	(C)	(D)
Data				
σ_X	2.10	1.25	6.65	1.08
$Corr(X_t, X_{t-1})$	0.68	0.56	0.54	0.71
$Corr(X_t, GDP_t)$	1	0.88	0.91	0.81
Baseline Vintage				
σ_X	0.56	0.32	3.60	0.57
$Corr(X_t, X_{t-1})$	0.89	0.88	0.81	0.80
$Corr(X_t, GDP_t)$	1	0.24	0.88	0.84
Neoclassical model $\overline{\xi_l}, \overline{\xi_h} = 0$				
σ_X	0.33	0.24	3.24	0.49
$Corr(X_t, X_{t-1})$	0.57	0.53	0.41	0.40
$Corr(X_t, GDP_t)$	1	-0.39	0.90	0.88

Note: Statistics computed over a simulation of 5000 periods with HP-filtered series obtained with a penalty of 100.

In Appendix E, we show that the result of technology adoption as a source of amplification of business cycle dynamics is not a byproduct of extracting the cyclical component by applying the HP-filter. Specifically, in Table *A*.5, we report business cycle statistics computed on the stationary variables of the model (as both models share the same deterministic trend).

7.2 Inspecting the Mechanism: Impulse Response Functions

In this section, we discuss the propagation mechanism of technology shocks in the vintage model and in the neoclassical benchmark. We compute impulse response functions by tracing the model dynamics following a unitary standard deviation shock to the efficiency of the latest vintage ($q_{0,t}$) when the economy sits at its stochastic steady state which is computed by simulating the model forward assuming that $q_{0,t}$ is constant at its mean value.

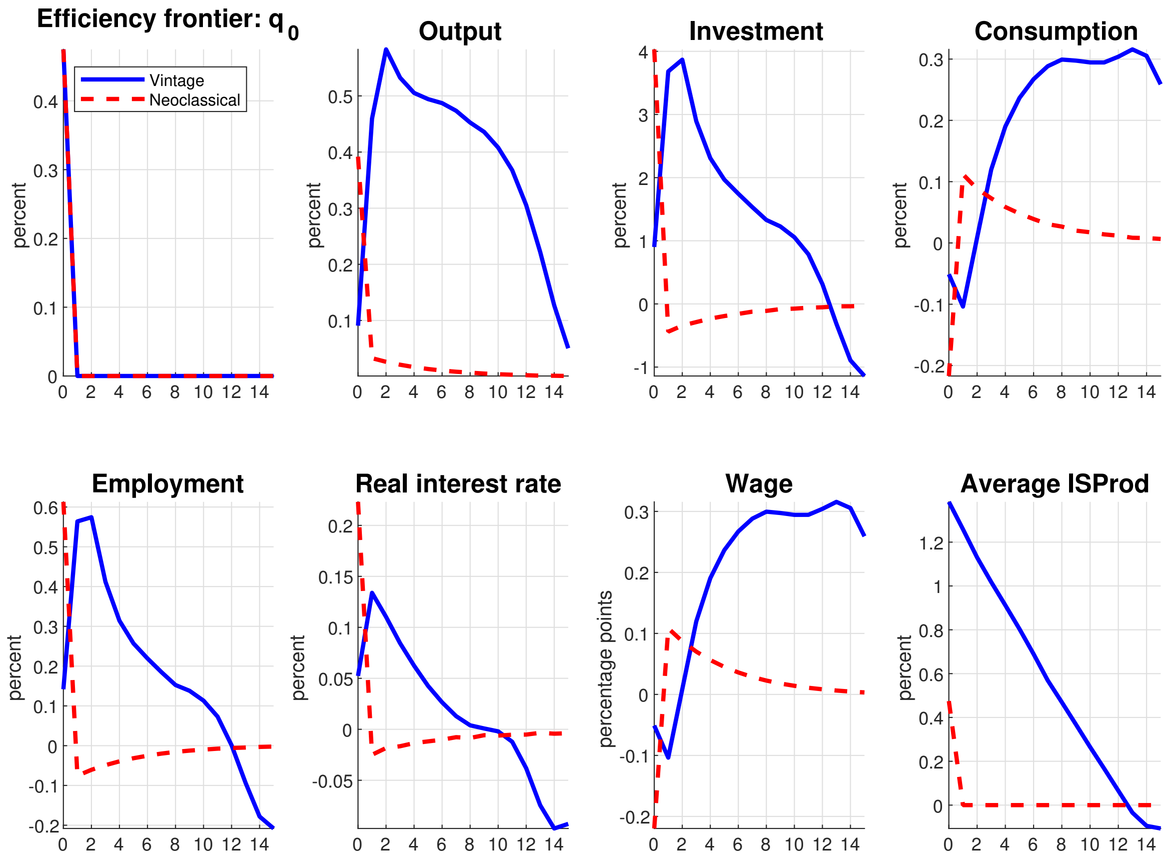

Figure 3: Impulse response functions following an increase with zero persistence in the efficiency of the latest vintage q_0 in the vintage model (solid line) and in the neoclassical benchmark with $\bar{\xi}_{l,h} = 0$ (dashed line).

Figure 3 contrasts the dynamics of the neoclassical benchmark with the vintage model. As in Den Haan, Ramey and Watson (2000), to highlight the differences in the propagation mechanism, we focus on the zero-persistence case where the temporary increase in the latest vintage lasts only one period. In the neoclassical model (dashed red lines), firms increase investment and employment as the latest technology becomes more efficient. Simultaneously, households finance higher investments by cutting consumption to enjoy a higher real interest rate. The zero persistence in q_0 implies that the dynamics of the aggregate series are short-lived as they revert back to the stochastic steady state immediately after the shock. In the vintage model (blue continuous lines), the magnitude of the aggregate fluctuations is larger relative to the neoclassical model as the technology adoption motives make investment, productivity, and employment more sensitive to the shock. The aggregate shock makes the current vintage available to firms more obsolete by increasing the technology distance relative to the frontier. The positive technology shock increases the benefit of switching to the latest and more productive vintage of investment goods, prompting firms to pay the nonconvex cost. The resulting boost to investment-specific productivity leads to an increase in economywide efficiency. This higher efficiency supports the prolonged increase in investment, employment, and output, well beyond the initial increase in the efficiency of the latest vintage. The average investment-specific productivity ("Average ISProd") increases much more in the vintage model and beyond the initial input of the exogenous shock as firms invest in newer, more productive investment goods. By strengthening the propagation mechanism of investment-specific technology shocks, the persistence of aggregate series in the vintage model is higher than in the benchmark neoclassical model and closer to the data.

Our results emphasize the role of firm investment heterogeneity in shaping aggregate dynamics; a finding counter with the "irrelevance" result of lumpy investment obtained in Thomas (2002) and Khan and Thomas (2003, 2008). In those papers, general equilibrium effects, namely, adjustment in the real interest rate and the wage, dampened the increase in investment demand due to nonconvex capital adjustment costs. The key difference in our framework that overcomes the irrelevance result is the productivity gains associated with switching to newer, more efficient investment goods. In models based on Solow (1960), such as Khan and Thomas (2003), firms can only acquire the latest vintage of the investment good. In our framework, firms optimally choose whether to obtain the newest vintage or keep acquiring the current vintage. The efficiency gap between the latest and the current vintage implies that adopting the latest vintage results in efficiency gains above the one implied by the stochastic evolution of the frontier. As discussed above, this adoption channel is critical for amplifying aggregate dynamics. Moreover, as shown by the amplified responses of the real interest rate and the wage, general equilibrium effects are at play in our framework, dampening the response of aggregate variables relative to partial equilibrium settings.

8 Concluding Remarks

We study the role of microeconomic heterogeneity for business cycle dynamics. Using firm-level data, we show that investment age, the time elapsed between large investment episodes, proxies for the technology operated by firms. Motivated by this empirical evidence, we formulate a general equilibrium model with rich firm heterogeneity where latest technology are embodied in the newest investment goods as in Solow (1960) and Greenwood, Hercowitz and Krusell (1997, 2000). In the model, a non-convex adjustment cost limits the firm's ability to switch from less efficient to newer more efficient investment goods. A positive technology shock induces more firms to pay the adjustment cost and to invest in more efficient capital goods. As a result, aggregate prouctivity increases boosting GDP, investment, and employment.

As technology is embodied in new capital goods, microeconomic heterogeneity in the efficiency of investment goods strengthens the propagation mechanism of technology shocks, amplifying, for a given persistence and size of the shock, the magnitude of aggregate fluctuations relative to the neoclassical growth model. Our results support the view that microeconomic heterogeneity is relevant for the propagation of business cycle dynamics.

References

Bachmann, Rüdiger, and Christian Bayer. 2014. "Investment Dispersion and the Business Cycle." *American Economic Review*, 104(4): 1392–1416.

Bachmann, Rüdiger, Ricardo J. Caballero, and Eduardo M. R. A. Engel. 2013. "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model." *American Economic Journal: Macroeconomics*, 5(4): 29–67.

Becker, Randy A., John Haltiwanger, Ron S. Jarmin, Shawn D. Klimek, and Daniel J. Wilson. 2006. "Micro and Macro Data Integration: The Case of Capital." *A New Architecture for the U.S. National Accounts*, 541–610.

Bertolotti, Fabio, Alessandro Gavazza, and Andrea Lanteri. 2023. "Dynamics of Expenditures on Durable Goods: The Role of New-Product Quality." *The Economic Journal*, 133(652): 1641–1656.

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

Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. 2018. "Really Uncertain Business Cycles." *Econometrica*, 86(3): 1031–1065.

Boucekkine, Raouf, David de la Croix, and Omar Licandro. 2011. "Vintage Capital Growth Theory: Three Breakthroughs." *UFAE and IAE Working Papers, mimeo*, 875.11.

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

Christiano, Lawrence J., and Jonas D. M. Fisher. 1998. "Stock Market and Investment Good Prices: Implications of Macroeconomics." *Federal Reserve Bank of Chicago*, , (WP-98-6).

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

Cooper, Russell W., and John C. Haltiwanger. 2006. "On the Nature of Capital Adjustment Costs." *Review of Economic Studies*, 73(3): 611–633.

Cummins, Jason G., and Giovanni L. Violante. 2002. "Investment-Specific Technical Change in the US (1947-2000): Measurement and Macroeconomic Consequences." *Review of Economic Dynamics*, 5(2): 243–284.

Den Haan, Wouter 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.

Den Haan, Wouter J., Garey Ramey, and Joel Watson. 2000. "Job Destruction and Propagation of Shocks." *American Economic Review*, 90(3): 482–498.

Doms, Mark E., and Timothy Dunne. 1998. "Capital Adjustment Patterns in Manufacturing Plants." *Review of Economic Dynamics*, 1(2): 409–429.

Fiori, Giuseppe. 2012. "Lumpiness, Capital Adjustment Costs and Investment Dynamics." *Journal of Monetary Economics*, 59(4): 381–392.

Fiori, Giuseppe, and Filippo Scoccianti. 2023. "The Economic Effects of Firm-Level Uncertainty: Evidence Using Subjective Expectations." *Journal of Monetary Economics*, 140: 92–105.

Gordon, Robert J. 1990. *The Measurement of Durable Goods Prices. NBER Books*, NBER, Inc.

Gourio, Francois, and Anil K Kashyap. 2007. "Investment Spikes: New Facts and a General Equilibrium Exploration." *Journal of Monetary Economics*, 54(Supplement): 1–22.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell. 1997. "Long-Run Implications of Investment-Specific Technological Change." *American Economic Review*, 87(3): 342–362.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell. 2000. "The Role of Investment- Specific Technological Change in the Business Cycle." *European Economic Review*, 44(1): 91–115.

Hansen, Gary D. 1985. "Indivisible Labor and the Business Cycle." *Journal of Monetary Economics*, 16(3): 309–327.

Hopenhayn, Hugo A. 1992. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica*, 60(5): 1127–1150.

House, Christopher L. 2014. "Fixed Costs and Long-Lived Investments." *Journal of Monetary Economics*, 68: 86–100.

Hulten, Charles R. 1992. "Technical Change is Embodied in Capital." *IMF Economic Review*, 82(4): 964–980.

Johansen, Leif. 1959. "Substitution versus Fixed Production Coefficients in the Theory of Economic Growth." *Econometrica*, 29: 157–176.

Judd, Kenneth L. 1998. *Numerical Methods in Economics.* MIT Press.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2010. "Investment shocks and business cycles." *Journal of Monetary Economics*, 57(2): 132–145.

Khan, Aubhik, and B. Ravikumar. 2002. "Costly Technology Adoption and Capital Accumulation." *Review of Economic Dynamics*, 5(2): 489–502.

Khan, Aubhik, and Julia K. Thomas. 2003. "Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities Matter?" *Journal of Monetary Economics*, 50(2): 331–360.

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

Khan, Aubhik, and Julia K. Thomas. 2013. "Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity." *Journal of Political Economy*, 121(6): 1055–1107.

Knittel, Christopher R. 2011. "Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector." *American Economic Review*, 101(7): 3368–3399.

Krusell, Per, and Anthony A. Smith. 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy*, 106(5): 867–896.

Licandro, Omar, Maria Reyes Maroto Illera, and Luis Puch. 2005. "Innovation, Machine Replacement and Productivity." *C.E.P.R. Discussion Papers*, 1(5422).

Øivind, Nilsen A., and Fabio Schiantarelli. 2003. "Zeros and Lumps in Investment: Empirical Evidence on Irreversibilities and Nonconvexities." *The Review of Economics and Statistics*, 85(4): 1021–1037.

Power, Laura. 1998. "The Missing Link: Technology, Investment, and Productivity." *The Review of Economics and Statistics*, 80(2): 300–313.

Sakellaris, Plutarchos. 2004. "Patterns of Plant Adjustment." *Journal of Monetary Economics*, 51(2): 425–450.

Sakellaris, Plutarchos, and Daniel J. Wilson. 2004. "Quantifying Embodied Technological Change." *Review of Economic Dynamics*, 7(1): 1–26.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review*, 97(3): 586–606.

Solow, Robert. 1960. "Investment and Technological Progress." *Mathematical Methods in Social Sciences 1959, 89-104*, , ed. S. Karlin K. Arrow and P. Supper, Chapter 7. Stanford University Press.

Tauchen, George. 1986. "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions." *Economics Letters*, 20(2): 177–181.

Thomas, Julia K. 2002. "Is Lumpy Investment Relevant for the Business Cycle?" *Journal of Political Economy*, 110(3): 508–534.

Winberry, Thomas. 2021. "Lumpy Investment, Business Cycles, and Stimulus Policy." *American Economic Review*, 111(1): 364–96.

Wolff, Edward N. 1996. "The Productivity Slowdown: The Culprit at Last? Follow-Up on Hulten and Wolff." *American Economic Review*, 86(5): 1239–1252.

Online Appendix to "Aggregate Dynamics and Microeconomic Heterogeneity: The Role of Vintage Technology"

A Data Sources

Detailed information on yearly balance sheets comes from Cerved Group S.P.A. (Cerved database), while data on employment and wages are obtained from the Italian National Institute of Social Security (INPS). Industry-specific price deflators and depreciation rates are obtained from the Italian National Statistical Institute (ISTAT). Sectors are constructed by aggregating available data from two-digit industries, according to the 2007 NACE classification. The agricultural sector includes industries 1, 2, 3, and 8. The manufacturing sector comprises industries 10, 11, and 13-33.

Table A.1: Sectoral Data

Sector	No. of Obs.
Agriculture, forestry, and fishing	96,087
Manufacturing	1,487,826
Electricity and gas supply	12,324
Water supply	40,249
Construction	614,258
Wholesale and retail trade	1,324,078
Transportation and storage activities	189,789
Accommodation and food service	267,581
Information and communication	223,826
Financial and insurance activities	25,160
Real estate activities	60,759
Professional, scientific, and technical activities	224,766
Administrative and support service activities	172,656
Public administration and defense	31,138
Education	121,044
Human health and social work	66,950
Other activities	46,403

Electricity and gas supply includes industry 35. The water supply sector includes industries 36-39. The construction sector includes industries 41-43. The wholesale and retail trade sector includes industries 45-47. The transportation and storage activities sector includes industries 49-53. The accommodation and food service sector includes industries 55 and 56. The information and communication sector includes industries 58-63. The financial and insurance activities sector includes industry 66. The real estate activities sector includes industry 68. The professional, scientific, and technical activities sector includes industries 69-75. The administrative and support service activities sector includes industries 77-82. The public administration and defense sector includes industry 85. The education sector includes industries 86-88. The human health and social work sector includes industries 90-93. The other activities sector includes industries 95 and 96. The composition of the data set by sector is reported in Table $A.1$. Our data on expected sales growth comes from the Survey of Industrial and Service Firms (INVIND), a large annual business survey conducted by the Bank of Italy on a representative sample of firms. Since 2002, the reference universe in INVIND consists of firms with at least 20 employees operating in industrial sectors (manufacturing, energy, and extractive industries) and nonfinancial private services, with administrative headquarters in Italy. The survey adopts a one-stage stratified sample design. The strata are combinations of the branch of activity (according to an 11-sector classification), size class (in terms of number of employees classified in seven buckets), and region in which the firm's head office is located. In recent years, each wave has had around 4,000 firms (3,000 industrial firms and 1,000 service firms). The data are collected by the Bank of Italy's local branches between February and April every year. The advantage of INVIND, relative to Cerved, is that it provides managers' expectations about future sales. The data set has a panel dimension. The firms observed in the previous edition of the survey are always contacted again if they are still part of the target population. In contrast, those no longer wishing to participate are replaced with others in the same branch of activity and size class. To limit the impact of outliers, we winsorize the 1% tails of the average expected sales.

B Investment Rates and Total Factor Productivity

Our measure of interest is TFP, together with investment age. Next, we discuss the construction of intermediate variables. Our computations follow the prevalent practice in the existing literature.

B.1 Total Factor Productivity

As in Bloom et al. (2018), we measure value-added $v_{f,t}$ for each firm f at year t as

$$v_{f,t} = Q_{f,t} - M_{f,t}, \tag{A.1}$$

where $Q_{f,t}$ is nominal output and $M_{f,t}$ is cost of materials. Nominal quantities are deflated by the corresponding sectoral deflators to obtain a measure of real value-added. Concerning labor input, we directly observe the wage bill and the number of employees for the firm at a given time t. We follow Bloom et al. (2018) and define value-added-based TFP as

$$log(\hat{z}_{f,t}) = log(v_{f,t}) - \theta_f log(k_{f,t}) - \nu_f log(N_{f,t}), \tag{A.2}$$

where $v_{f,t}$ denotes real value added, $k_{f,t}$ the real capital stock, $N_{f,t}$ labor input, and θ and ν are the cost shares for capital and labor, respectively. We follow Bachmann and Bayer (2014) and estimate θ and ν by the median of the firm average share of factor expenditure in total value-added, as defined by

$$\begin{aligned}\hat{\theta}_f &= T^{-1}\sum_t \frac{wn_{f,t}}{v_{f,t}} \text{ and}\\ \hat{\nu}_f &= T^{-1}\sum_t \frac{(r_{f,t}+\delta_{f,t})k_{f,t}}{v_{f,t}},\end{aligned} \tag{A.3}$$

where $wn_{f,t}$ is the real wage bill and $r_{f,t}$ the real cost of funds for the corporate sector and is estimated using the average real interest rate on banking loans for the corporate sector. As in Becker et al. (2006) and most of the existing literature, we construct the real capital stock series using the perpetual inventory method so that

$$k_{f,t} = (1-\delta_{f,t})k_{f,t-1} + i_{f,t}, \tag{A.4}$$

where $i_{f,t}$ is real net investment (deflated using sectoral deflators for capital expenditures) on tangible and intangible assets. To initialize the recursion, we estimate the real stock of capital using the book value of fixed assets net of funds amortization. The depreciation rate δ is common within sectors.

C Equilibrium and (S,s) Decision Rules

Given the presence of fixed cost, the adoption and investment decisions are akin to exercising an option. Consider a firm of a type $(\varepsilon, z, \tilde{k})$ drawing adjustment cost ξ. Define the value associated with the value of action V^A and the one with the inaction choice V^{NA} as

$$V^A(\varepsilon_t, q_{0,t}, \tilde{k}^A_{t+1}; q_{0,t}, \mu_t) \equiv \max_{\tilde{k}^A_{t+1}\in\mathbf{R}_+} R(\varepsilon_t, q_{0,t}, \tilde{k}^A_{t+1}; q_{0,t}, \mu_t), \tag{A.5}$$

$$V^{NA}(\varepsilon_t, q_{j,t}, \tilde{k}^{NA}_{t+1}; q_{0,t}, \mu_t) \equiv \max_{\tilde{k}^{NA}_{t+1}\in\mathbf{R}_+} R(\varepsilon_t, q_{j,t}, \tilde{k}^{NA}_{t+1}; q_{0,t}, \mu_t), \tag{A.6}$$

Next, define the firm's target capital $\tilde{k}^A_{t+1}$ as the optimal choice of $\tilde{k}$—when the firm obtains capital embodying the latest vintage that solves the right-hand side of (A.5). The solution to the problem in (A.5) is independent of the current stock of capital $\tilde{k}_t$ and ξ_t, but not ε_t (and of course $q_{0,t}$), given persistence in firm-specific productivity. As a result, all firms with current productivity ε_t and pay their fixed costs to obtain capital embodying the latest vintage of technology choose a common target capital for the next period, $\tilde{k}^A_{t+1} = k(\varepsilon_t, q_{0,t}, \mu_t)$, and achieve a common gross value V^A. By contrast, firms that do not pay adjustment costs have value V^{NA}. In this case, the firm keeps its current vintage $q_{j,t}$, which becomes more obsolete (i.e., more distant from the technological frontier) at a rate γ_q. The firm also adjusts its stock of capital consistent with its current vintage and idiosyncratic productivity.

A firm will pay the fixed cost if $V^A - w\xi$—the value of adjusting—is at least as great as V^{NA} — the value of inaction. Given continuity in the adjustment cost ξ, it is possible to identify threshold value such that a type $(\varepsilon_t, q_{j,t}, \tilde{k}_t)$ firm is indifferent between action and inaction:

$$-w(q_{0,t},\mu_t)\hat{\xi}(\varepsilon_t,q_{j,t},\tilde{k}_t;q_{0,t},\mu_t)+V^A(\varepsilon_t,q_{0,t},\tilde{k}_t;q_{0,t},\mu_t)=V^{NA}(\varepsilon_t,q_{j,t},\tilde{k}_t;q_{0,t},\mu_t). \qquad (A.7)$$

To summarize the adoption and investment decision define $\xi^T(\varepsilon_t,q_{j,t},\tilde{k}_t;q_{0,t},\mu_t)\equiv\min[\bar{\xi},\hat{\xi}(\varepsilon_t,q_{j,t},\tilde{k}_t;q_{0,t},\mu)]$ so that $0\le\xi^T(\varepsilon_t,q_{j,t},\tilde{k}_t;q_{0,t},\mu_t)\le\bar{\xi}$. Any firm $(\varepsilon_t,q_{j,t},\tilde{k}_t)$ that draws an adjustment cost at or below its type-specific threshold ξ^T will pay the fixed cost and acquire capital embodying the latest vintage $q_{0,t}$.

Thus, for a given group of firms of type $(\varepsilon_t,q_{j,t},\tilde{k}_t;q_{0,t},\mu_t)$, a fraction $G\left[\xi^T(\varepsilon_t,q_{j,t},\tilde{k}_t;q_{0,t},\mu_t)\right]$ pay their fixed cost to adopt the latest vintage and optimally choose capital. Thus, the market-clearing levels of consumption required to determine $p(q_{0,t},\mu_t$ using equation 8 is given by

$$\begin{aligned}C \;=\; & \int_S \varepsilon F(\tilde{k},n) & (A.8)\\ & -\int_S G\left[\xi^T(\varepsilon_t,q_{j,t},\tilde{k}_t;q_{0,t},\mu_t)\right]J(\xi_t\le\xi^T(\varepsilon_t,,q_{j,t},\tilde{k}_t;q_{0,t},\mu_t))i^A \\ & -\int_S\left[1-G\left[\xi^T(\varepsilon_t,q_{j,t},\tilde{k}_t;q_{0,t},\mu_t)\right]\right]J(\xi_t>\xi^T(\varepsilon_t,,q_{j,t},\tilde{k}_t;q_{0,t},\mu_t))i^{NA},\end{aligned}$$

where it is understood that i^A and i^{NA} depend upon the firm's current state. Finally, we turn to the evolution of the firm distribution, $\mu_{t+1}=\Gamma(q_{0,t},\mu_t)$. It is useful to define the indicator function $J(x)=1$ if x is true, and 0 otherwise. For each $(\varepsilon_t,q_{j,t},\tilde{k}_t)\in S$

$$\begin{aligned}&\mu_{t+1}(\varepsilon_{t+1},q_{j,t+1},\tilde{k}_{t+1})\\ =\;&\sum_{l=1}^{N_\varepsilon}\pi^\varepsilon_{lm}\left[\begin{array}{c}\int J(\xi_t\le\xi^T(\varepsilon_t,q_{j,t},\tilde{k}_t;q_{0,t},\mu_t))G\left[\xi^T(\varepsilon_l,q_{j,t},\tilde{k}_t;q_{0,t},\mu_t)\right]\mu_t(\varepsilon_l,q_{j,t},\tilde{k}_t)\\ +\int J(\xi_t>\xi^T(\varepsilon_t,,q_{j,t},\tilde{k}_t;q_{0,t},\mu_t))\left[1-G\left[\xi^T(\varepsilon_l,q_{j,t},\tilde{k}_t;q_{0,t},\mu_t)\right]\mu_t(\varepsilon_l,q_{j,t},\tilde{k}_t)\right]\end{array}\right]\end{aligned}$$

D Computation

D.1 Solution Algorithm

To solve the model, we employ an extension of the method in Krusell and Smith (1998) and Khan and Thomas (2008, 2013). This strategy replaces the aggregate law of motion for the distribution with a forecast rule. Typically, to predict prices and the future proxy aggregate state, agents use the mean capital stock. In our framework, the distribution is defined over capital stocks and the vintage technologies chosen by the firms. To obtain accurate forecast rules, we use the mean of the capital stocks, the mean of investment-specific productivity, and their interaction. This set of regressors works very well, yielding an accurate forecast of prices and future proxy aggregate state. Given the perceived laws of motion we solve the individual firm's problem by value iteration and obtain the policy functions. We then simulate the model for 5,000 periods, and update the price functions and the perceived law of motions by explicitly imposing market clearing. Specifically, the value function to solve the firm's problem defined in equation 4, 5, and 6 are the basis of our numerical solution of the economy. The solution algorithm involves repeated application of the contraction mapping implied by equation 4, 5, and 6 to solve for firms' value function, given the price functions $p(\mu, q_0)$. More specifically, the firm's problem amounts to find the next-period value of capital $\tilde{k}_{t+1}$. To do so, we use a golden section search to allow for continuous control. The process for the idiosyncratic process ε and q_0 are approximated using the procedure in Tauchen (1986) over 7 and 5 possible, respectively. Given the growth rate γ_{q_0}, we discretize the process for q_0 spanning 1.2 standard deviations above and below mean, to ensure that there is no technology regress. This choice ensures that, in presence of zero adjustment cost, firms in the neoclassical model finds optimal to adjust technology even with the lowest realization of the aggregate shock. We compute the value function exactly at the grid points above and interpolate for in-between values. This procedure is implemented using a multi-dimensional cubic splines procedure, with a so-called "not-a-knot"-condition to address the large number of degrees of freedom problem, when using splines (see Judd, 1998). With the firm's policy function at hand, we compute the stationary distribution and verify that the guessed price is consistent with market clearing. We update the guessed price function $p(\mu, q_0)$ until convergence, i.e., until the guessed and the market-clearing price converges.

Table A.2: Forecasting Rules - Vintage Model

	Technology	Max Error	Mean Error	$S.E.$	R^2
	(A)	(B)	(C)	(D)	(E)
Forecasting $\log(\overline{K}_{t+1})$					
	$q_{0,1}$	0.00329	0.00064	0.00088	0.99877
	$q_{0,2}$	0.00379	0.00087	0.00114	0.99780
	$q_{0,3}$	0.00420	0.00091	0.00118	0.99783
	$q_{0,4}$	0.00419	0.00089	0.00117	0.99773
	$q_{0,5}$	0.00331	0.00072	0.00093	0.99820
Forecasting $\log(\overline{Q}_{t+1})$					
	$q_{0,1}$	0.00331	0.00065	0.00085	0.99485
	$q_{0,2}$	0.00288	0.00068	0.00087	0.99553
	$q_{0,3}$	0.00214	0.00036	0.00047	0.99909
	$q_{0,4}$	0.00392	0.00087	0.00101	0.99330
	$q_{0,5}$	0.00375	0.00099	0.00121	0.98110
Forecasting $\log(p_t)$					
	$q_{0,1}$	0.00127	0.00029	0.00038	0.99387
	$q_{0,2}$	0.00125	0.00030	0.00039	0.99359
	$q_{0,3}$	0.00128	0.00029	0.00037	0.99437
	$q_{0,4}$	0.00121	0.00034	0.00042	0.99331
	$q_{0,5}$	0.00146	0.00032	0.00041	0.99389

We estimate a regression conditional on each realization of the aggregate process, $q_{0,t}$. In Tables *A*.2 and *A*.3, we assess the accuracy of the forecasting rule for both models. For simplicity, we report only one set of results for the neoclassical benchmark model as the accuracy is virtually the same.

Table *A*.3: Forecasting Rules - *RBC* Model

	Technology	Max Error	Mean Error	*S.E.*	R^2
	(A)	(B)	(C)	(D)	(E)
Forecasting $\log(\overline{K}_{t+1})$					
	$q_{0,1}$	0.00003	0.00002	0.00003	0.99999
	$q_{0,2}$	0.00003	0.00002	0.00003	0.99999
	$q_{0,3}$	0.00004	0.00002	0.00003	0.99999
	$q_{0,4}$	0.00004	0.00002	0.00003	0.99999
	$q_{0,5}$	0.00002	0.00001	0.00003	0.99999
Forecasting $\log(p_t)$					
	$q_{0,1}$	0.00000	0.00002	0.00001	0.99999
	$q_{0,2}$	0.00001	0.00003	0.00001	0.99999
	$q_{0,3}$	0.00001	0.00002	0.00001	0.99999
	$q_{0,4}$	0.00001	0.00002	0.00001	0.99999
	$q_{0,5}$	0.00001	0.00002	0.00001	0.99999

We find that the algorithm yields a very accurate solution as reflected in the high R^2 and small standard errors. As discussed by Den Haan (2010), the values of R^2 are averages and scaled by the variance of the dependent variable, so the lower variance of $\overline{Q}_{t+1}$ relative to the other approximated variables explains its somewhat lower R^2 even for a similar mean and maximum errors. To provide a robust statistic, we perform a long simulation of the model (5,000 periods) and compare the paths of the approximated variables with an alternative simulation obtained iterating only on the estimated law of motions. We compute the maximum and mean distance of each approximated variable from the value taken in the actual simulation. These values are reported in Table *A*.4. As all variables are in logs, the errors can be interpreted as percentage deviations. The results show that the solution is not only accurate in the sense that it produces accurate one-step ahead forecast, but also in that the forecast errors do not accumulate over time.

Table $A.4$: Forecasting Rules - Den Haan Test

	$\log(\overline{K}_{t+1})$	$\log(\overline{Q}_{t+1})$	$\log(p_t)$
	(A)	(B)	(C)
Baseline Vintage			
Maximum Error	0.00983	0.00859	0.00341
Mean Absolute Error	0.00194	0.00171	0.00070
Neoclassical model $\bar{\bar{\xi}}_{l,h} = 0$			
Maximum Error	0.00265	n.a.	0.00041
Mean Absolute Error	0.00102	n.a.	0.00031

Notes: Maximum and mean errors reported in the table are obtained in simulations of the vintage model based on subjective laws of motion.

E Unfiltered Business Cycle Moments

Here, we report the business cycle moments of the economy without applying the HP-filter for the vintage model and its neoclassical counterparts. The goal of this exercise is to highlight that the amplification result is not a product of filtering the data.

Table $A.5$: Unfiltered Business Cycle Statistics: Technology Shock

	GDP	Consumption	Investment	TFP	Employment
	(A)	(B)	(C)	(D)	(E)
Baseline Vintage					
σ_X	0.82	0.49	4.44	0.31	0.69
$Corr(X_t, X_{t-1})$	0.94	0.94	0.87	0.92	0.86
$Corr(X_t, GDP_t)$	1	0.54	0.86	0.90	0.80
Neoclassical model $\bar{\xi}_{l,h} = 0$					
σ_X	0.42	0.29	3.53	0.08	0.61
$Corr(X_t, X_{t-1})$	0.71	0.67	0.50	0.61	0.49
$Corr(X_t, GDP_t)$	1	-0.06	0.86	0.70	0.83

Note: Statistics report the business cycle moments of stationary variables in each of the models considered.