Aggregate Dynamics and Microeconomic Heterogeneity: The Role of Vintage Technology*

Giuseppe Fiori Filippo Scoccianti

February 13, 2024

Abstract

We study the role of firm heterogeneity in vintage technology for business cycle dynamics in a general equilibrium model. In the data, we document that investment age—the time elapsed since the firm experienced a significant investment episode—proxies for the obsolescence of the technology operated by firms and is negatively correlated with firms' total factor productivity (TFP). Motivated by this evidence, we incorporate a technology adoption decision in an otherwise neoclassical growth model with firm heterogeneity. A non-convex adoption cost prevents firms from adopting the most recent technology every period. As new and old vintages coexist in equilibrium, the non-degenerate distribution of capital stocks and technologies across firms determines aggregate TFP. Idiosyncratic and aggregate shocks that alter the timing of technology adoption at the firm level shift the cross-sectional distribution of TFP and generate endogenous procyclical movements in economy-wide productivity beyond the ones accounted for by exogenous shocks. Through this vintage technology adoption channel, microeconomic heterogeneity in investment age amplifies the magnitude of business cycle fluctuations relative to a benchmark neoclassical growth model.

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

*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) and Khan and Thomas (2008). Since the work of Johansen (1959) and Solow (1960), a large body of the theoretical literature has emphasized the link between technology adoption and investment, because the newer vintages are of better quality and may enhance the efficiency of existing capital.

This paper studies the role of investment, vintage technology, and total factor productivity (TFP) heterogeneity in 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 TFP. We then study the aggregate implications of heterogeneity in investment age and formulate a state-of-the-art model that builds on Khan and Thomas (2008) in which a firm's TFP results from a technology adoption choice and exogenous factors. In the model, a non-convex adoption cost prevents firms from adopting the most recent technology every period. New and old vintages coexist in equilibrium, yielding a non-degenerate distribution of capital stocks and technologies across firms. Idiosyncratic and aggregate shocks alter the timing of technology adoption at the firm level and result in shifts in the cross-sectional distribution of TFP. This shift generates endogenous procyclical movements in economy-wide productivity beyond the ones accounted for by exogenous shocks. Through this vintage technology channel, microeconomic heterogeneity in investment age *amplifies* the magnitude of business cycle fluctuations relative to a benchmark neoclassical growth model.

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 and in line with many studies across advanced economies, investment at the firm level is a large and infrequent, or *lumpy*, episode.[1] On average, 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; managers' expectations about future sales; and instrumenting investment age with differences between planned and actual investment and the reasons behind it. 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 (2008) and features a quality ladder structure in the form of vintage technology adoption and frictionless capital choice. In the model, the firm's productivity includes *(i)* the *permanent* vintage component, which is endogenous to the timing of technology adoption, and *(ii)* a temporary but persistent idiosyncratic component, which is entirely exogenous. The vintage component evolves as in a quality ladder model, with firms catching up to a frontier that grows stochastically. Firms optimally decide if and when to adopt the latest vintage—i.e., the 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 adopt the latest technology, while others postpone it. Conditional on the adoption decision, firms optimally choose the capital stock. In equilibrium, technologies of different quality coexist. The history of adoption choices contributes to microeconomic

[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.

heterogeneity in TFP and, in turn, determines aggregate productivity. Instead, when the adoption cost is set 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 where firms' and aggregate productivity are exogenous.

Our framework accounts well for the microeconomic heterogeneity in investment and TFP and salient features of business cycle dynamics. On the microeconomic side, the model is consistent with the cross-sectional distributions of investment rates and investment age. Also, the model accounts for about half 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 model also reproduces the near-zero autocorrelation of investment rates, the persistence in firm-level TFP, and the magnitude of the empirical relationship between investment age and TFP estimated using firm-level data. On the macroeconomic side, the macroeconomic series' magnitude and comovement are in line with the data when the source of aggregate fluctuations are technology shocks —i.e. shocks to the efficiency of the latest vintage.

The critical result of the analysis is that the cross-sectional distribution of investment age, which proxies for vintage technology, is relevant for business cycle dynamics. 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 TFP. The endogenous response of productivity constitutes an additional force that amplifies macroeconomic dynamics relative to a neoclassical growth model. For a given persistence, the vintage technology model requires shocks about 40 per cent smaller than the neoclassical growth model to generate business cycles of the same magnitude. We validate this mechanism against the data by showing that the model reproduces the relationship between TFP and investment age at the *sectoral* level. Our results support the view that microeconomic heterogeneity is relevant for aggregate dynamics, and prolonged investment slumps can contribute to stagnant 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 technology adoption, 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); Greenwood, Hercowitz and Krusell (1997); and Greenwood, Hercowitz and Krusell (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) can be interpreted as a measure of investment-specific technological progress. There is, instead, 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.[2] We extend the existing studies by showing that invest-

[2]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

ment 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), 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 if and when to adopt the latest vintage and, conditional on this choice, the next-period stock of capital. Significantly, vintage technology increases TFP dispersion beyond the one implied by exogenous productivity. We reproduce similar evidence on investment heterogeneity relying on technology adoption costs rather than capital adjustment costs. Our framework replicates several critical moments of the microeconomic structure of the economy, for instance, the cross-sectional distribution of investment rates, the empirical relationship between investment age and TFP, and the time-series properties of firm-level investment rates and TFP.

The literature has also debated on the relevance of accounting for the cross-sectional dynamics in investment for aggregate dynamics, see Thomas (2002), House (2014), Fiori (2012), Khan and Thomas (2008), Bachmann, Caballero and Engel (2013), and recently Winberry (2021). Our results indicate that technology adoption motives constitute a power amplification mechanism of aggregate disturbances, providing a channel through which microeconomic heterogeneity shapes business cycle dynamics.

The main difference with models that study vintage capital based on Solow (1960) is that, to be consistent with the empirical analysis, we focus on the total effect of embodied technical change on firms' TFP. The quantitative focus of our analysis distinguishes our work

recently (and using Spanish manufacturing data) Licandro, Maroto Illera and Puch (2005) find the opposite result.

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.[3] For a complete list of references and a historical perspective on the evolution of the literature on vintage capital, see the extensive surveys of Boucekkine, de la Croix and Licandro (2011) and Boucekkine and de Oliveira Cruz (2015).

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 source 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

[3] Cooley, Greenwood and Yorukoglu (1997) study the balanced growth path and the transitional dynamics of a deterministic model with two sectors and vintage capital and compare it with the neoclassical growth model. Samaniego (2006) formulates a model that emphasizes the role of organizational capital as friction that prevents firms from adopting newer technology.

Social Security (INPS) (see Appendix 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 2015, 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 the trade sectors constitute more than one-half of the observations in the data.[4]

INVIND is 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 sample extends from 1996 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 and on the difference between planned and actual investment that we exploit in the empirical analysis to establish a robust predictive 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 1998-2015 sample. 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

[4] In Table A.1, we report the composition of the data set by sector. 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.

of disinvestment.Investment $I_{f,t}$ includes expenditures on equipment and structures as they 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 Output (B)	Share of Investment (C)	Share of Employment (D)
$ik \geq 20\%$	20.04%	26.77%	61.04%	27.52%
$-5\% \leq ik \leq 5\%$	37.21%	25.67%	5.76%	27.01%
$ik \leq -20\%$	2.32%	1.98%	-6.65%	2.14%

Note: ik denotes the investment rate. See the main text for the definition. The distribution of investment rates is computed over the sample period 1998 to 2015.

that exhibit a spike—i.e., an investment rate above 20 percent. On average, these investors account for 61 percent 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] While the share of investment that these two groups of firms account for differs substantially, the share of output and employment are equivalent.

2.3 Investment Age Proxies 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.

[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.

Investment age, denoted by $Inv.Age_{f,t}$, is based on the time elapsed since the last 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 an investment spike, the variable $Inv.Age_{f,t}$ equals zero, progressively increasing by one yearly until the same firm experiences an investment spike. The long time-series dimension in our data allows us to split the sample using the first half to initialize the distribution of $Inv.Age_{f,t}$ and the second half of the sample in the empirical analysis.

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 estimating a negative correlation: Firms with higher investment age employ less advanced. 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 evidence supports the interpretation of $Inv.Age_{f,t}$ as a firm-level measure of the gap 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

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

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

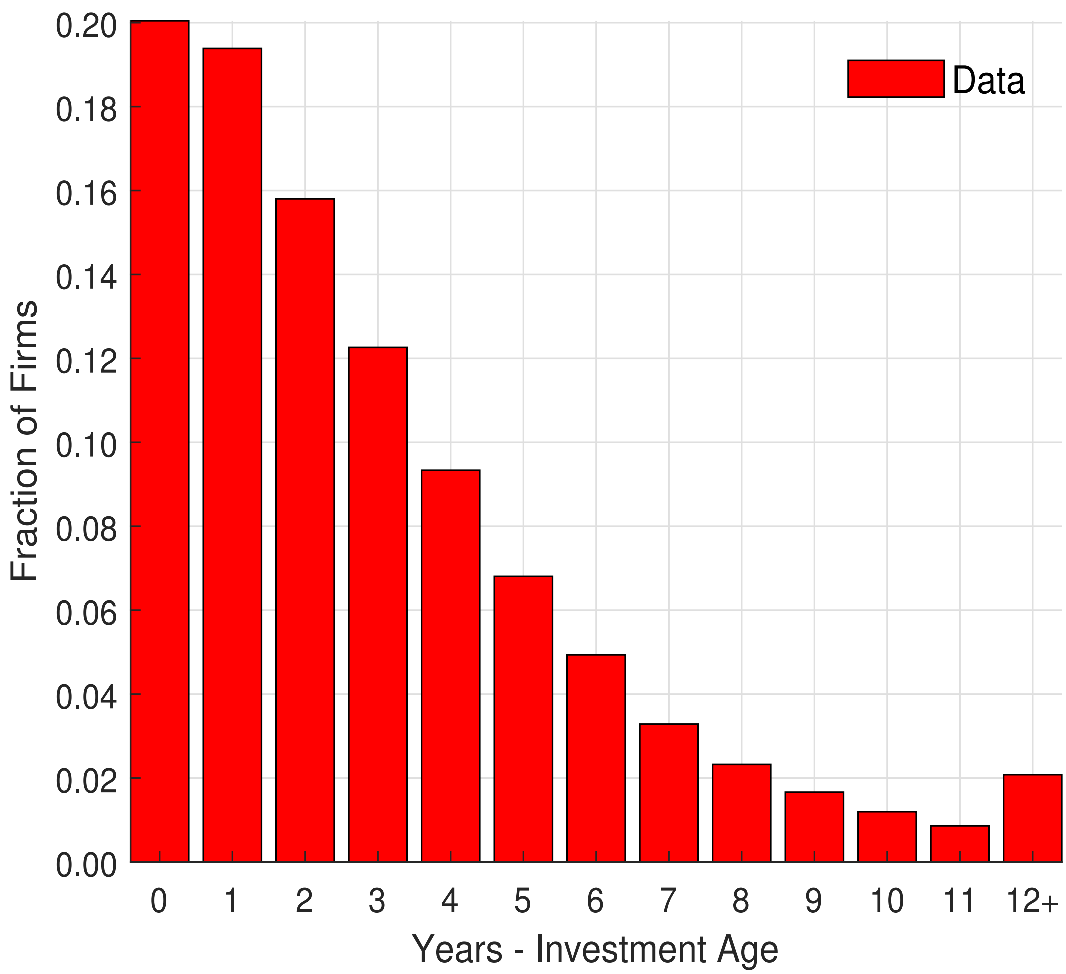

The average firm's investment cycle occurs every three and a half years. Intuitively, average investment age is negatively correlated with GDP growth. During a GDP expansion, as in the early 2000s, the share of firms that experienced an investment spike increased to about 22 per cent. During GDP contractions, 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. In other words, 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. This empirical relationship supports 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 a panel of Italian firms:

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

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 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.

The set of controls features sector, time, and sector-time dummies to account for unobserved time-variant and time-invariant industry-specific characteristics and aggregate factors, 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).

One key element of our analysis is the inclusion of managers' expectations about the one-year-ahead growth rate of sales elicited in INVIND 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.

To further clean out the empirical relationship, we employ an instrumental variable (IV) approach exploiting firms' self-reported discrepancies between actual and planned investments. INVIND elicits firms to report when i) their actual and planned investments differ and ii) the reasons behind this discrepancies. When there is a discrepancy, the indicator takes value one and zero otherwise. Note that we do not observe the magnitude of the gap and whether the gap between actual and planned investments is positive or negative. These indicators are instruments to estimate equation 1.

We instrument $Inv.Age_{f,t-1}$ with three dummy variables that take a value of one when firms' planned and actual investments differ because of different-than-expected purchasing prices, legislative changes, delays in delivering capital goods, and the second lag of investment age. Intuitively, we assume that $Inv.Age_{f,t-1}$ follows as an autoregressive process of order one, where we instrument spikes that reset $Inv.Age_{f,t-1}$ using the discrepancy between planned and actual investment. In our sample, about 10 percent of the firms report investment being different from what they had planned.

To be clear, despite our efforts in dealing with the potential endogeneity of investment age, we prefer interpreting our estimates in a *predictive* rather than a causal sense, given the plausible violation of the exclusion restriction assumption.

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, by postponing large capital adjustments, firms delay the introduction of new technologies in the production process, suffering lower TFP as the gap between their current technology 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 estimation methods. 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, quantifying the productivity loss at about 0.9 percent per year. Column B shows that entry dynamics do not drive results. Focusing on firms that are 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 in INVIND to include managers' expectations about future sales, $s^e_{avg,f,t}$, to the set of controls in our baseline specification. Column D indicates that the role of investment age for TFP dynamics survives controlling for firms' expectations. In addition, news about the future proxied by firms' expected sales predicts higher TFP.[8]

IV estimates, shown in columns E and F, confirm delaying investment prevents firms from upgrading their technology and results in a lower TFP. The strength of the F-statistics above conventional values indicates a well-identified first stage. Reassuringly, the Sargan overidentifying test does not reject the validity of our instruments with p-values of 0.65 and 0.46. Concerning the magnitude of coefficients, point estimates in columns E and F are somewhat larger than OLS estimates, but their difference is not statistically significant at

[8] In Fiori and Scoccianti (2023), we 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)	$TFP_{f,t}$ (E)	$TFP_{f,t}$ (F)
$Inv.Age_{f,t-1}$	-0.860^{***}	-0.837^{***}	-0.930^{***}	-0.906^{**}	-1.333^{**}	-1.214^{**}
	(0.00)	(0.00)	(0.01)	(0.08)	(0.04)	
$s^{e}_{avg,f,t}$				0.370^{***}		0.341^{***}
				(0.02)		(0.00)
N. of obs.	2,756,422	2,596,223	3,347	2,773	2,098	1,625
R^2	0.746	0.750	0.919	0.918		
Anderson Stat					121.67	139.85
Cragg-Donald Stat					425.76	422.63
Sargan Stat: p-val					0.65	0.46
Estimator	OLS	OLS	OLS	OLS	IV	IV
Firm FE	✓	✓	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Time×Ind. FE	✓	✓	✓	✓	✓	✓
Source of Data	Cerved		Cerved matched with INVIND			
Sample Period	1998-2016	1998-2016	2003-2016	2003-2016	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 E and F report estimates obtained instrumenting $Inv.Age_{f,t}$, see the text for details on instruments.

conventional values.[9]

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. Motivated by this evidence, we now formulate an equilibrium model that accounts for the documented firm-level evidence to study the aggregate implications of technology adoption.

[9]The 95 percent confidence interval around the IV coefficient for $Inv.Age_{f,t-1}$ is $[-2.493, -0.173]$ and contains the OLS estimate. We cannot reject the null hypothesis of the estimated coefficient on $Inv.Age_{f,t-1}$ in column D because it is not statistically different from its IV counterparts in column F at a 10 percent level.

4 Model: Firm Heterogeneity and Technology Adoption

We formulate a general equilibrium framework featuring firm heterogeneity and technology adoption to study the aggregate implications of investment age heterogeneity. Our starting point is the neoclassical growth model of Khan and Thomas (2008), the benchmark for quantitative analysis involving firm dynamics in a general equilibrium framework. The main innovation is endogenizing firms' TFP by introducing a vintage technology structure. The firm's productivity depends not only on exogenous idiosyncratic factors but also on the technological vintage chosen by the firm. In other words, vintage technology is a defining characteristic of the firm. The firm's problem involves deciding the optimal timing to obtain the latest vintage and how much to invest. This option is subject to a non-convex adjustment cost that, in equilibrium, leads to the coexistence of vintages of different quality and contributes to reproducing the firm-level evidence on capital accumulation.

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

4.1 Production

The economy consists of a continuum of firms normalized to one and features one commodity that can be consumed or invested. Each firm has access to an increasing and concave production function that combines a predetermined capital stock k and labor hired on a spot market with its available technology to produce output y:

$$y = \varepsilon z k^{\theta} n^{\nu}, \tag{2}$$

where $0 < \theta + \nu < 1$.[10] Production efficiency depends upon two variables, ε and z. ε denotes the idiosyncratic productivity that is exogenous to the firm. z identifies the current vintage of technology available to each production unit and is optimally chosen by the firm. Every period, firms decide whether to pay the cost ξ and adopt the latest vintage or keep the current technology vintage and postpone adoption. The technological frontier grows stochastically at the gross rate of $\gamma_A > 1$. Along the balanced growth path, z_0 indicates the latest vintage or the technological frontier that, deflated by its trend, is equal to 1. A firm that chooses not to obtain the latest vintage keeps its current technology that becomes more obsolete relative to the frontier at a per-period rate γ_A, so that $z' = z/\gamma_A$. In the next section, we describe the tradeoffs associated with the choice of technology adoption. As in Khan and Thomas (2008), $\varepsilon \in \{\varepsilon_1, \varepsilon_2, ..., \varepsilon_{N_\varepsilon}\}$, 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 $l = 1, ..., N_\varepsilon$. Similarly, $\gamma_A \in \{\gamma_{A,1}, \gamma_{A,2}, ..., \gamma_{A,N_{\gamma_A}}\}$ where $Pr(\gamma_A = \gamma_{A,q} | \gamma_A = \gamma_{A,n}) \equiv \pi^{\gamma_A}_{nq} \geq 0$, and $\sum_{q=1}^{N_{\gamma_A}} \pi^{\gamma_A}_{nq} = 1$ for each $n = 1, ..., N_{\gamma_A}$. In each period, a firm is defined by its vintage productivity z, its idiosyncratic productivity level $\varepsilon \in \mathcal{E} \equiv \{\varepsilon_1, \varepsilon_2, ..., \varepsilon_{N_\varepsilon}\}$, its predetermined stock of capital $k \in \mathbf{R}_+$, and its cost associated with vintage adoption $\xi \in [0, \overline{\xi}]$, which is denominated in units of labor. In every period, the plant chooses its current level of employment; production occurs; labour is paid the real wage, denoted by w.

4.2 Firm's Technology Adoption and Investment Decision

The technology adoption decision is subject to a cost, while capital is chosen freely. In every period, each firm chooses between keeping its current vintage or adopting the latest technology. This choice involves deciding whether to pay its current adoption cost ξ. By paying ξ, the firm obtains the latest vintage z_0 and optimally chooses the stock of capital k', where primes denote next-period variables. The firm's capital stock evolves according to

[10] Variables reported are deflated by their respective stochastic trends. Along the balanced growth path, γ_A denotes the gross trend growth rate of the technology frontier. Consumption and capital grow at a gross rate $\gamma = \gamma_A^{1/(1-\theta)}$.

$\gamma_k k' = (1-\delta)k + i$, where i is its current investment and $\delta \in (0,1)$ is the rate of physical capital depreciation.[11] Implicitly, we assume full retrofitting of the capital stock, which is that the productivity associated with the new vintage applies to the capital stock of the firm already installed.[12] Specifically, by forfeiting ξ units of current labor, the firm can invest in any future capital $k \in \mathbf{R}_+$ and upgrade technology z to the latest vintage of technology z_0.

Firms that postpone paying the adjustment cost keep their current vintage and can undertake investment i as capital can be freely adjusted. In this case, the firm's distance from the technological frontier or the degree of technological obsolescence increases so that $z' = z/\gamma_A$, and $k' \in \mathbf{R}_+$.

As in Khan and Ravikumar (2002), the adoption adjustment cost ξ is non-convex, and its modelling 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. As in Thomas (2002), we assume that ξ is independently and identically distributed across firms and time. Every period, each firm draws its current cost of vintage adoption $\xi \geq 0$ (denominated in units of labor) from the time-invariant distribution G common to all production units. As the firm's current adjustment cost has no implication for its future adjustment, the distribution of firms is summarized by (ε, z, k) : the idiosyncratic productivity ε, vintage technology z, and capital stock k. To characterize the distribution of firms over (ε, z, k), 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 (γ_A, μ), the growth rate of the technology frontier and the distribution of firms

[11] As in Solow (1960), the dynamics of vintage models cannot be represented through a representative firm unless knife-edge conditions are met. This case is invalid in models with a vintage structure and non-convex adjustment costs. Motivated by our empirical evidence in Section 3, our identifying assumption is that the firm vintage, linked to the capital accumulation decision of the firm, is a defining characteristic of the firm. Khan and Thomas (2003) employ a different identifying assumption. They assume that the latest vintage technology applies only to the most recent investment.

[12] Alternatively, one could assume that adopting the new technology comes at the cost of scrapping a fraction of the firm's current k. This choice would amount to rescaling the fixed adoption cost to make it firm-specific.

that evolves according to a mapping (defined below) $\Gamma : \mu' = \Gamma(\gamma_A, \mu)$.

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 in 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 $p = p(\gamma_A, \mu)$. This variable indicates the pricing kernel used by firms to price output streams. Given the i.i.d. nature of the adjustment cost ξ, continuation values can be integrated out of future continuation values.

Let $v^1(\varepsilon_l, z, k, \xi; \gamma_A, \mu)$ denote the expected discounted value of a firm entering the period with (ε_l, z, k) and drawing an adjustment cost ξ when the aggregate state of the economy is $(\gamma_A,\ \mu)$. The dynamic optimization problem for the typical firm is described using a functional equation defined by equations 3, 4, and 5. 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 (ε_l, z, k):

$$V^0(\varepsilon_l, z, k; \gamma_A, \mu) = \int_0^{\bar{\xi}} V^1(\varepsilon_l, z, k, \xi; \gamma_A, \mu) dG(\xi). \tag{3}$$

The firm's profit-maximization problem, which takes as given the evolution of the firm distribution, $\mu' = \Gamma(\gamma_A, \mu)$, is then described by

$$V^1(\varepsilon, z, k, \xi; \gamma_A, \mu) \quad = \quad \max_{k^A,\, k^{NA}} \left\{ \begin{array}{c} [F(\varepsilon, z, k) - w(\gamma_A, \mu)L + (1-\delta)k]\, p(\gamma_A, \mu) + \\ \max \left[\begin{array}{c} -\xi w(\gamma_A, \mu) + R(\varepsilon, z_0, k^A; \gamma_A', \mu'), \\ R(\varepsilon, z/\gamma_A, k^I; \gamma_A', \mu') \end{array} \right] \end{array} \right\} \tag{4}$$

$s.t.\ k^A \in \mathbf{R}_+$ and $k^I \in \mathbf{R}_+$,

where $R(\varepsilon, z', k'; \gamma_A', \mu')$ represents the continuation value associated with a given combination of the idiosyncratic shock, the vintage, and the stock of capital:

$$R(\varepsilon, z, k^{'}; \gamma_A^{'}, \mu^{'}) \equiv -\gamma k^{'} p(\gamma_A, \mu) + \beta \sum_{q=1}^{N_{\gamma_A}} \sum_{m=1}^{N_\varepsilon} \pi_{nq}^{\gamma_A} \pi_{lm}^{\varepsilon} V^0(\varepsilon_m, z^{'}, k^{'}; \gamma_A^{'}, \mu^{'}) \tag{5}$$

Every period, the firm decides whether to pay the fixed cost ξ, upgrade its vintage, and adjust its capital stock accordingly. Otherwise, the firm keeps its current vintage. 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 $Z(\varepsilon, z, k, \xi; \gamma_A, \mu)$ and $K(\varepsilon, z, k, \xi; \gamma_A, \mu)$ represent the choice of technology and capital for the next period by firms of type (ε, z, k) with adjustment cost ξ.

4.4 Households

The economy features a continuum of identical households with access to a complete set of state-contingent claims. As there is no heterogeneity across households, these assets are in zero net supply in equilibrium. Moreover, they own shares in the production units, denoted by the measure $\lambda(\varepsilon, z, k; \gamma_A, \mu)$ and value $\rho_0(\varepsilon, z, k; \gamma_A, \mu)$. Given the value for their current shares, $\rho_0(\varepsilon, z, k; \gamma_A, \mu)$ and the real wage obtained for labor provided w, households maximize their lifetime expected utility by choosing current consumption c, labor effort L, as well as the numbers of $\lambda(\varepsilon, z, k)$ to purchase at prices $\rho_1(\varepsilon, z, k; \gamma_A, \mu)$:

$$W(\lambda; \gamma_A, \mu) = \max_{c, L, \lambda^{'}} \left[U(c, L) + \beta \sum_{q=1}^{N_{\gamma_A}} \pi_{nq} W(\lambda^{'}; \gamma_A^{'}, \mu^{'}) \right] \tag{6}$$

subject to

$$
\begin{aligned}
& c+\int_S \rho_1(\varepsilon, z, k;\gamma_A,\mu)\lambda^{'}\left(d\left[\varepsilon^{'}\times z^{'}\times k^{'}\right]\right) \\
\leq\ & w(\gamma_A,\mu)L\int \rho_0(\varepsilon, z, k;\gamma_A,\mu)\lambda\left(d\left[\varepsilon\times z\times k\right]\right).
\end{aligned} \tag{7}
$$

The household optimality condition yields:

$$
U_C w(\gamma_A,\mu)=U_L\left(c,L\right) \text{ and } p(\gamma_A,\mu)=U_C\left(c,L\right). \tag{8}
$$

Let us denote the optimal choices for the household as $C(\lambda;\gamma_A,\mu)$, $L^h(\lambda;\gamma_A,\mu)$, and $\Lambda^h(\varepsilon^{'},z^{'},\lambda,k^{'};\gamma_A,\mu)$.

4.5 Recursive Equilibrium

A recursive competitive equilibrium is a set of functions (p, v^1, Z, K, W, C, Λ^h, Γ) that satisfy the firms' and households' problem and clear the markets for assets, labor, and output:

(i) Firm's optimality: Taking p as given, $V^1\left(\varepsilon,z,k,\xi;p\right)$ solves equations 3, 4, and 5 and the corresponding policy functions $Z=Z\left(\varepsilon,z,k,\xi;p\right)$ and $K=K\left(\varepsilon,z,k,\xi;p\right)$.

(ii) Household's optimality: Taking p as given, the household's decisions satisfy equations in 8 and (C,L^h,Λ^h).

(iii) $\Lambda^h(\varepsilon_m,z,k;\gamma_A,\mu)=\mu(\varepsilon_m,z,k)$ for each $(\varepsilon_m,z,k)\in S$.

(iv) Commodity market clearing: $C=\int y d\mu-\int\int_0^{\overline{\xi}}\left[\gamma K\left(\varepsilon,z,k,\xi;\gamma_A,p\right)-\left(1-\delta\right)K\right]dGd\mu$ $-\int \xi dGd\mu$.

(v) Labor market clearing: $L^h=\int N\left(\varepsilon,z,k;\gamma_A,\mu\right)d\mu+\int_0^{\overline{\xi}}\xi J(x)dGd\mu$, where $J(x)=0$, if the plant does not upgrade its vintage, and 1 otherwise.

(vi) Model-consistent dynamics: The evolution of the cross-sectional distribution that characterizes the economy, $\mu'=\Gamma(\gamma_A,\mu)$, 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 Aggregate TFP

In this section, we discuss the role of the vintage structure for aggregate TFP. Non-convex adoption cost implies that the firm's technology adoption decision follows an (S,s) rule: Some firms adopt the latest vintage, while others postpone it.[13] Conditioning on the adoption decision, firms decide on next-period capital stock. The degree of obsolescence of the vintage currently available to the firm (the productivity gap from the technological frontier) and the realization of the idiosyncratic and aggregate shocks affect the timing of technology adoption at the firm level. Thus, in equilibrium, vintages of different qualities coexist. The history of adoption choices contributes to TFP heterogeneity and determines the *economy-wide* production efficiency. Shifts in the cross-sectional distribution, determined by variations in the firms' adoption decision and driven by shocks, result in aggregate TFP fluctuations that constitute an additional force in the propagation of aggregate shocks shaping investment and output dynamics.

The evidence in Sections 2 and 3 highlights the relationship between investment age and TFP: Firms with higher investment age (measured as the time elapsed between investment spikes) display, other things being equal, lower TFP than firms with lower investment age. This link is not hard-wired into the model. Adopting the latest vintage entails a non-convex adoption cost independent of the investment size necessary to reach the target capital.

Before describing the calibration strategy of the model, it is worth examining the role of

[13] See Appendix C for additional details.

three key parameters. The upper support of the adjustment cost distribution ($\bar{\xi}$) determines the magnitude of the cost to adjust technology. The higher $\bar{\xi}$, the higher the potential cost of adopting the latest vintage. Increasing this parameter leads to a higher average investment age. 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 first- and second-order moments of the Italian economy. Table 3 summarizes parameter values, targeted moments, and data sources. We are to assign values to 12 parameters related to the growth rate of aggregate variables along the balanced growth path in the absence of shocks ($\overline{\gamma}$ and $\overline{\gamma_A}$), the production process (δ, θ, and ν), individual preferences on disutility of labor and discount factor (A and β), the technology adjustment cost function ($\bar{\xi}$), the exogenous idiosyncratic productivity process (ρ_ε, and σ_ε), and the evolution of the technological frontier (ρ_{γ_A}, and σ_{γ_A}). We first describe the set of parameters that are externally calibrated—i.e., using independent evidence. Then, we focus on those estimated within the model to reproduce relevant targets in the data.

Externally Calibrated. 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 to 0.18 and 0.64, respectively, the median of the sector-level estimates employed to construct TFP in Section 3.1. We set the average output growth rate ($\overline{\gamma}$) to 0.8 percent, its sample average computed using data from 1988-2016.

Internally Calibrated. 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 to 1.23 to reproduce employment equal to 0.6. Given $\overline{\gamma}$ and θ, balanced growth rate restrictions pin down the mean growth rate of the technology frontier ($\overline{\gamma_A}$) equal to 0.65 percent. The discount factor β is set to 0.975 to reproduce the real annual interest rate in the data.

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
Disutlity of labor	A	1.234	Employment rate = 60%
Mean growth rate of output	$\bar{\gamma}$	1.008	Data
Mean growth rate of technology frontier	$\overline{\gamma_A}$	1.007	Balanced growth path restriction
Discount factor	β	0.975	Annual real interest rate = 2.3%
Persistence exogenous idiosyncratic productivity	ρ_ε	0.352	TFP persistence in the data
St. dev. idiosyncratic productivity	σ_ε	0.066	Share of firms experiencing $ik > 0.20$
Upper support of the adoption cost distribution	$\bar{\xi}$	113.229	Average $Inv.Age$
Persistence aggregate technology	ρ_{γ_A}	0.150	Data
St. dev. aggregate technology	σ_{γ_A}	0.003	ik distribution

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 upper support of the technology adjustment cost function ($\bar{\xi}$) is set to reproduce the average investment age of the firms' cross-sectional distribution and is equal to 113.2 implying high adjustment costs equal to about 5 per cent of aggregate investment. The persistence of the exogenous idiosyncratic productivity process (ρ_ε) is set so that the model firm's TFP displays the same persistence as the one estimated in the data (0.87). This coefficient obtains fitting an autoregressive process of order one to the logarithm of firm-level TFP without any firm-specific or time effects. 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 against a standard neoclassical benchmark with firm heterogeneity. Our vintage model nests the neoclassical benchmark as a special case that obtains when the technology adjustment cost, $\bar{\xi}$, is equal to zero, and there is no technological regress. Finally, the log of $\gamma_{A,t}$ follows an autoregressive process centered around its mean value $\overline{\gamma_A}$: $log(\gamma_{A,t}) = (1-\rho_{\gamma_A})\, log(\overline{\gamma_A}) + \rho_{\gamma_A} log(\gamma_{A,t-1}) + \sigma_{\gamma A}\varepsilon_{\gamma_A}$, where ε_{γ_A} is a normally distributed i.i.d. process with standard deviation $\sigma_{\gamma A}$. In the absence of empirical guidance for the evolution of the technological frontier, we choose ρ_{γ_A} so that the autocorrelation of γ_A matches the persistence of TFP in the data and σ_{γ_A} to reproduce the volatility of the growth rate of aggregate investment in the data for the 1988-2016 sample, consistent with the sample of the empirical analysis. Specifically, then set ρ_{γ_A} equal to 0.15 and $\sigma_{\gamma A}$ to 0.0026.

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] In theory, this could complicate the solution algorithm by requiring agents to forecast the behavior of two distributions rather than one. In practice, especially when the persistence of the aggregate shock is relatively low, the standard rule

[14] We note that, even when two firms have the same TFP, they may still choose different stock of capital next period as the continuation value depends upon ε and z.

that uses the mean of the capital stocks as a regressor works very well, yielding an accurate forecast of prices and future proxy aggregate state. Alternatively, one could forecast the mean of the cross-sectional distribution of TFP. Due to the high correlation between the two distributions of TFP and capital stocks, this choice is immaterial. Thus, we forecast the mean capital K' and the marginal utility of consumption p using $log(K')=\beta_0 + \beta_1 log(K) + \varepsilon_K$ and $log(p)=\beta_0 + \beta_1 log(K) + \varepsilon_p$. Moreover, we estimate the rule conditional on each realization of the aggregate process, $\gamma_{A,t}$. We report additional details regarding the solution method in Appendix D.

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 Sections 6.3 and 6.4, we show that the model accounts for a significant fraction of the dispersion in idiosyncratic TFP in the data and the empirical relationship at the firm level between TFP and investment age.

6.1 Cross-Sectional Distribution of Investment Rates

We now examine the model performance starting from the cross-sectional distribution of investment rates. As in Cooper and Haltiwanger (2006) and Khan and Thomas (2008), the cross-sectional distribution is summarized using five groups: inaction, positive and negative investment, 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	Positive Spikes	Negative Spikes	Positive Investment	Negative Investment
	$\lvert ik \rvert \leq 0.05$	$ik > 0.20$	$ik < -0.20$	$ik > 0.05$	$ik < -0.05$
	(A)	(B)	(C)	(D)	(E)
Data	37.21%	20.04%	2.32%	57.95%	4.80%
Baseline Vintage	22.87%	20.04%	3.52%	55.61%	21.52%
Neoclassical $\bar{\xi} = 0$	20.99%	26.04%	2.32%	65.02%	13.99%

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

The average share of firms experiencing a 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. Instead, the model overstates the share of firms downsizing capital and understates the share of inactive firms. A comparison with a neoclassical benchmark that retains the same set of parameters as the vintage model and sets the adjustment cost to technology, $\bar{\xi}$ equal to zero, sheds light on the role of vintage technology in driving aggregate dynamics. Relative to the neoclassical benchmark ($\bar{\xi} = 0$), 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 gap 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.

We also investigate the time-series implications of firm-level investment rates in the data and the model as an additional moment against which to validate our framework. 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 0.00 autocorrelation, close to 0.10 in the data.

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 that has elapsed since the last time the firm experienced an investment spike. The cross-sectional distribution of investment age displays a sizable and negative correlation (-0.61) with the cross-sectional distribution of TFP and, thus, conveys information about vintage technology.

We remind the reader that our calibration targets the average investment age of 3.4 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 of investment age relative to the data. Overall, the performance of the model is quite 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 investment age zero coincides by construction with the fraction of firms exhibiting spikes, the fraction of firms with age 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 $\overline{\xi}$ to zero, reduces the average investment age of the model to 1.5 years.

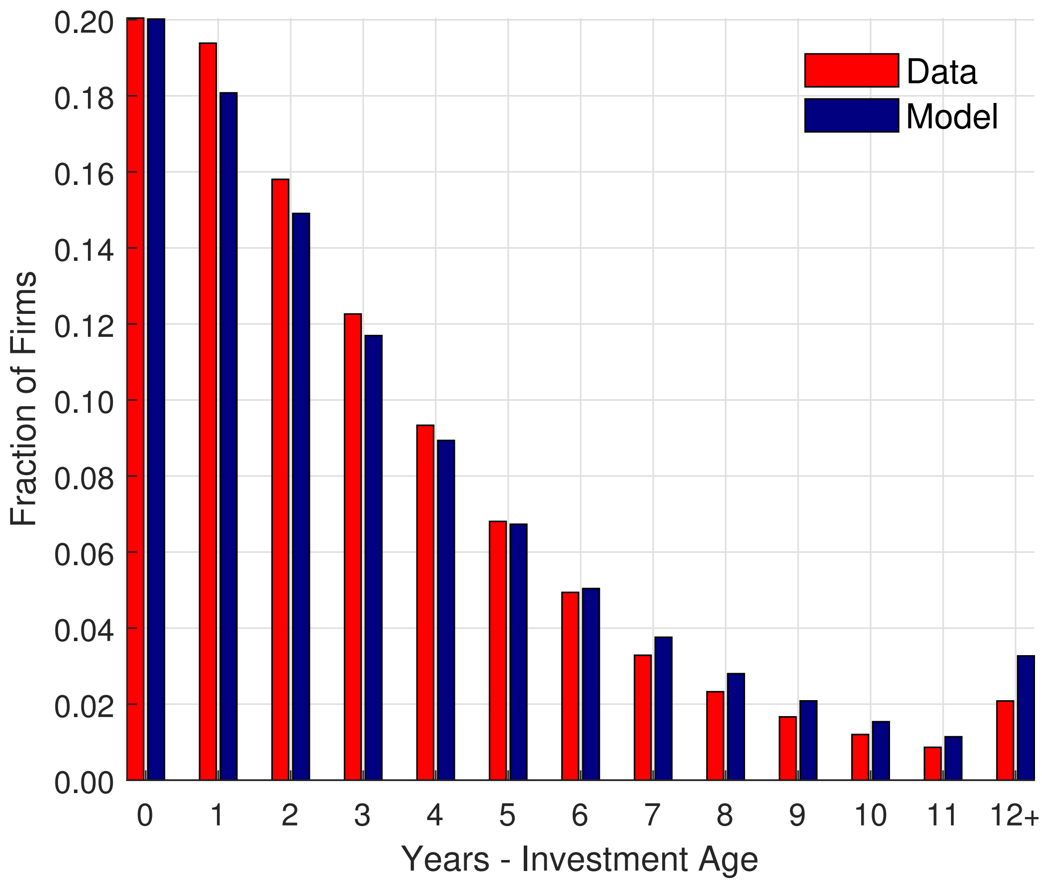

Figure 2: Comparison between the Empirical and the Model-Based Investment Age Distribution

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

We now discuss the time-series properties of firm-level TFP and its cross-sectional distribution in the vintage model. As mentioned in Section 5.1, the persistence of firm-level TFP is one of our calibration targets. In our model, the persistence of firm-level TFP is the by-product of two forces, the persistence of the vintage component z equal to 0.94 and the one of idiosyncratic exogenous component ε equal to 0.35.

Regarding the cross-sectional distribution of TFP, vintage technology amplifies heterogeneity in firm-level TFP relative to one implied by idiosyncratic factors ε. To assess the contribution of vintage technology to the cross-sectional dispersion, we compute the log of the interquartile and the 90/10 range of firm-level TFP in the model and compare it with the data. The latter obtains as dispersion of residuals from regressing TFP on sector, time, and sector-time effects over the 1998-2016 sample.

Table 5: Data and Model: TFP Dispersion

	Data	Baseline Vintage	Neoclassical $\overline{\xi} = 0$
	(A)	(B)	(C)
$\log(TFP_{75}/TFP_{25})$	0.511	0.211	0.081
$\log(TFP_{90}/TFP_{10})$	1.082	0.397	0.163
Source:	Data	Baseline Model	

Note: The $\log(TFP_{75}/TFP_{25})$ is the log of the ratio between the 75^{th} and 25^{th} (or 90^{th} and 10^{th}) percentiles of the cross-sectional distribution of TFP. The Neoclassical model retains the same parameters of the Baseline Vintage model described in Section 5.1 except for the upper support of the adjustment cost distribution $\overline{\xi}$ equal to zero.

The key result in Table 5 is that vintage technology nearly doubles the TFP dispersion relative to the neoclassical benchmark. As shown in columns B and C, the model's 75-25 and 90-10 ranges are about twice as much as the one in the data. In turn, the vintage model accounts for about 40 percent of the TFP dispersion in the data.

The vintage model generates a mildly negative correlation between TFP dispersion and output in deviations from the trend (-0.36). 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 technology adoption.

6.4 Model and Data: TFP and Investment Age

After establishing that our framework provides a good account of the cross-sectional distribution of investment rates, investment age, and firm-level TFP, we now use the model to interpret the empirical evidence between TFP and the lag of investment age in Section 3 using data simulated from the model.

This exercise is useful for two distinct reasons. It allows us to verify that the relationship between TFP and $InvAge$ in the model is in line with what is estimated in the data and to decompose the role of vintage effects from exogenous TFP fluctuations. In the model, TFP is the product of vintage effects, z, and exogenous idiosyncratic process, ε. Therefore, the coefficient on the log of $TFP_{f,t}$ is the sum of the coefficient on regressing z and ε on $Inv.Age_{f,t-1}$, respectively.[15]

Table 6: Data and Model: TFP and Investment Age

	$TFP_{f,t}$ Data	$TFP_{f,t}$ $(z \times \varepsilon)$	$TFP_{f,t}$ (z)	$TFP_{f,t}$ (ε)
	(A)	(B)	(C)	(D)
$Inv.Age_{f,t-1}$	-0.903	-0.855	-0.662	-0.193
Source:	Data	Baseline Model		

Note: The dependent variable is the log of total factor productivity (TFP) for firm f at time t. $Inv.Age_{f,t1}$ measures the time elapsed between investment spikes, defined as the firm experiencing an investment rate above 20 percent at time t-1.

Table 6 reports the results. Column A reports estimate in Section 3.2 to facilitate a comparison with the data. Turning to relationships based on simulated data, column B shows that the model displayed a negative relationship between the log of TFP and lagged $Inv.Age$.

[15] We simulate data for a sample of 30000 firms, enough points to replicate the cross-sectional distribution of investment rates, for 3000 periods. We take the latest 70 observations to estimate equation 1.

From a quantitative standpoint, the estimated coefficient of -0.85 is close in magnitude to the one in the data (-0.90). We note that the relationship between investment age and TFP is not hard-wired in the model; firms can adopt the latest vintage without exhibiting an investment spike. Given the focus on aggregate dynamics, we take comfort that the relationship between TFP and $Inv.Age$ closely matches the data, providing an additional moment against which to validate our framework.

In columns C and D, we decompose the role of vintage effects, z from idiosyncratic shock, ε, in driving the estimated relationship. The log of TFP is the sum of the logs of z and ε. Vintage effects (-0.66) rather than the idiosyncratic exogenous TFP process (-0.19) drive the estimated relationship between TFP and lagged $Inv.Age$.

As a placebo test, we run the same experiment setting the technology adjustment cost to zero, $\bar{\xi} = 0$, in a model that does not reproduce the cross-sectional distribution of investment age. There, the coefficient is negative (-0.22) but still far smaller than the one implied by vintage effects. The inclusion of firm-specific or time effects to the specifications using simulated data reduces the magnitude of the coefficients without changing the conclusions drawn from Table 6.

The framework with vintage effects provides a good account of the cross-sectional distributions of investment rates and age. It delivers the magnitude of the negative relationship between TFP and lagged $Inv.Age$.

7 Vintage Technology Amplifies Aggregate Fluctuations

In this section, we study the business cycle implications of heterogeneity in vintage technology. In Section 7.1, we show that the microeconomic heterogeneity induced by vintage technology amplifies the volatility of aggregate series in response to technology shocks relative to a neoclassical benchmark. In Section 7.2, we provide *industry-level* evidence on the empirical relationship between TFP and investment age that supports the mechanism embedded in the vintage model.

7.1 Business Cycle Properties

To assess the role of vintage technology in the business cycle, 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 $\overline{\xi} = 0$. Unlike the neoclassical growth model, in which firms are always at the technology 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.

While productivity and technology are often used interchangeably to label stochastic disturbances to production efficiency in the context of models belonging to the real business cycle literature, the same is not valid in our framework. As the technology frontier evolves stochastically, technology innovations directly affect only the adopters of the latest vintage, increasing the gap from the frontier for nonadopters. To assess quantitatively the amplification in aggregate dynamics due to technology adoption, we perform the conventional business cycle exercise by simulating the model in response to technology shocks.[16]

Table 7 reports business cycle statistics of growth rates of aggregate series in the data and counterparts for the baseline vintage and the neoclassical model.

The vintage model reproduces the investment volatility by construction and successfully reproduces the relative magnitude of the cyclical volatilities for consumption, investment, and TFP.

As shown in Table 7, vintage effects amplify aggregate dynamics relative to the nested

[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, $\gamma_{A,t}$, is always non-negative. Removing this assumption does not significantly affect the main results of the paper.

Table 7: Business Cycle Statistics Technology Shock

	ΔGDP	ΔConsumption	ΔInvestment	ΔTFP
	(A)	(B)	(C)	(D)
Data				
σ_X	1.965	1.532	5.346	1.405
$Corr(X, GDP)$	1	0.776	0.923	0.765
Baseline Vintage				
σ_X	0.808	0.221	5.346	0.387
$Corr(X, GDP)$	1	0.337	0.968	0.821
Neoclassical model $\bar{\xi} = 0$				
σ_X	0.630	0.182	3.014	0.272
$Corr(X, GDP)$	1	-0.044	0.975	0.902

Note: Each entry represents the volatility of the respective variable. Δ indicates the growth rate. C, I, and L refer to consumption, investment, and labor, respectively.

neoclassical benchmark. In response to the same shocks, the standard deviation of GDP in the vintage model increases by about 30 percent. The amplified response of TFP and investment explains this result. First, aggregate shocks alter the incentives to adopt new technologies, shifting the cross-sectional distribution of TFP and contributing to aggregate TFP fluctuations over and above the initial impulse of the stochastic technological frontier. Column D shows that aggregate TFP is more volatile relative to a neoclassical benchmark by about 50 percent. Second, a more volatile TFP amplifies the response of aggregate investment, contributing to the more volatile growth rate of GDP relative to the neoclassical benchmark. Here, we focus on the growth rates of aggregate series and show in Appendix E that the amplification of aggregate dynamics in the vintage model also holds when we report statistics for the cyclical component of each series extracted by applying the Hodrick and Prescott filter.

7.2 Industry-Level Evidence on TFP and Investment Age

In this section, we provide further support for the mechanism in the vintage model. To verify that the amplification mechanism in the model is consistent with the aggregate empirical relationship between TFP and investment age, we compute the semi-elasticity of TFP with respect to investment age in the model and compare it with the one estimated using industry data. On a panel of two-digit industries from the Cerved database over the 1998-2016 sample, we regress sectoral TFP on the lag of investment age, computed as average on the cross-section of firms, including sectoral fixed effects and year effects to control for potential aggregate trends and policy changes. In the model, the same specification is estimated using the log of TFP in deviations from the trend as a dependent variable.

Table 8: Aggregate Empirical Relationship between TFP and $Inv.Age$

	Data	Baseline vintage	Neoclassical $\bar{\xi} = 0$
	(A)	(B)	(C)
Semi-elasticity of TFP w.r.t. $Inv.Age_{sec.,t-1}$	-12.561^{**}	-4.768^{**}	n.a.
N. of obs.	1109	1109	
Sample period	1998-2016		

Note: Each entry represents the volatility of the respective variable. Δ indicates the growth rate. C, I, and L refer to consumption, investment, and labor, respectively.

Table 8 reports the results. We emphasize two main points. First, the data supports the negative empirical relationship between TFP and the lag of investment age at the industry level, validating the mechanism in the mode. In the neoclassical model, the empirical relationship is zero by construction. TFP is always on trend and does not vary in response to investment age. Second, the estimated relationship in the vintage model is statistically significant and comparable in magnitude with the one in the data.

8 Concluding Remarks

We study the role of microeconomic heterogeneity in vintage technology for business cycle dynamics. Using firm-level data, we show that investment age, the time elapsed between large investment episodes, proxies the technology operated by firms.

To study the aggregate implications of this microeconomic evidence, we formulate a general equilibrium model with rich firm heterogeneity characterized by a technology choice in which firms decide when to adopt the latest vintage that nests, as a special case, the neoclassical growth model with exogenous TFP heterogeneity. The model simultaneously accounts for microeconomic heterogeneity in investment and TFP, replicating the empirical relationship between TFP and investment age at the firm and aggregate level.

Because of the technology adoption mechanism, the cross-sectional distribution of TFP is relevant for aggregate dynamics and amplifies aggregate fluctuations relative to the neoclassical growth model. Our results support the view that a prolonged investment slump, like the one experienced by many advanced economies in the decade after the Great Recession, contributes to dampening aggregate productivity growth, raising questions about the recovery after the Covid-19 pandemic.

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.

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, and Bruno de Oliveira Cruz. 2015. "Technological Progress and Investment: A Non-Technical Survey." *AMSE Working Papers, mimeo*, 1519.

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.

Cooley, Thomas F., Jeremy Greenwood, and Mehmet Yorukoglu. 1997. "The Replacement Problem." *Journal of Monetary Economics*, 40(3): 457–499.

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.

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*, National Bureau of Economic Research, 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.

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.

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.

Samaniego, Roberto M. 2006. "Organizational Capital, Technology Adoption and the Productivity Slowdown." *Journal of Monetary Economics*, 53(7): 1555–1569.

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 aggregating available data from two-digit industries, according to the 2007 NACE classification. The agriculture 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

The 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 7 buckets), and region in which the firm's head office is located. In recent years, each wave has 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 to provide managers' expectations about future sales and indicators of investment delays that we will use as an instrument in our IV approach to measuring the effects of investment age on productivity. 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, and $N_{f,t}$ 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 decision is akin to exercising an option. Consider a firm of a type (ε, z, k) drawing adjustment cost ξ. Define the value associated with the value of action $V^A(\varepsilon, z, k; \gamma_A, \mu)$ and the one with the inaction choice $V^I(\varepsilon, z, k; \gamma_A, \mu)$ as

$$V^A(\varepsilon, z_0, k; \gamma_A, \mu) \equiv \max_{k' \in \mathbf{R}_+} R(\varepsilon, z_0, k'; \gamma_A, \mu), \tag{A.5}$$

$$V^I(\varepsilon, k; \gamma_A, \mu) \equiv \max_{k' \in \mathbf{R}_+} R(\varepsilon, z_0, k'; \gamma_A, \mu), \tag{A.6}$$

Next, define the firm's target capital k^* as the optimal choice of k—when the firm obtains 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 k and ξ, but not ε (and of course z_0), given persistence in firm-specific productivity. As a result, all firms with current productivity ε and pay their fixed costs to upgrade to their latest vintage choose a common target capital for the next period, $k^* = k(\varepsilon,\ \gamma_A,\ \mu)$, and achieve a common gross value $V^A(\varepsilon, z_0, k; \gamma_A, \mu)$. By contrast, firms that do not pay adjustment costs have value $V^I(\varepsilon, z, k; \gamma_A, \mu)$. In this case, the firm keeps its current vintage z that becomes more obsolete (i.e., more distant from the technological frontier) at a rate γ_A. The firm gets to adjust its stock of capital consistent with its current vintage and idiosyncratic shock.

A firm will pay the fixed cost if $V^A(\varepsilon, k; \gamma_A, \mu) - w(\xi)$—the value of adjusting—is at least as great as $V^I(\varepsilon, z, k; \gamma_A, \mu)$ — the value of inaction. Given continuity in the adjustment cost ξ, it is possible to identify threshold value such that a type (ε, z, k) firm is indifferent between action and inaction:

$$- w(\gamma_A, \mu)(\hat{\xi}(\varepsilon, z, k; \gamma_A, \mu)) + V^A(\varepsilon, k; \gamma_A, \mu) = V^I(\varepsilon, z, k; \gamma_A, \mu). \tag{A.7}$$

To summarize the adoption and investment decision define $\xi^T(\varepsilon, z, k; \gamma_A, \mu) \equiv \min [\ \overline{\xi}$, $\hat{\xi}$(ε,z,k;γ_A, μ)] so that $0 \leq \xi^T(\varepsilon, z, k; \gamma_A, \mu) \leq \overline{\xi}$. Any firm (ε, z, k) that draws an adjustment cost at or below its type-specific threshold $\xi^T(\varepsilon, z, k; \gamma_A, \mu)$ will pay the fixed cost and adjust k and z.

Thus, for a given group of firms of type $(\varepsilon, z, k; \gamma_A, \mu)$, a fraction $G\left[\xi^T(\varepsilon, z, k; \gamma_A, \mu)\right]$ pay their fixed cost to adopt the latest vintage and optimally choose capital. Thus, the market-clearing levels of consumption required to determine p using equation 8 is given by

$$\begin{aligned} C \;=\; & \int_S F(\varepsilon, z, k) && (A.8)\\ & -G\left[\xi^T(\varepsilon, z, k; \gamma_A, \mu)\right] J(\xi \leq \xi^T(\varepsilon, z, k; \gamma_A, \mu))(i^A - \xi) \\ & - \left[1 - G\left[\xi^T(\varepsilon, z, k; \gamma_A, \mu)\right]\right] J(\xi > \xi^T(\varepsilon, z, k; \gamma_A, \mu)) \left[i^I\, \mu\left(d\left[\varepsilon \times z \times k\right]\right)\right], && (A.9) \end{aligned}$$

where it is understood that i^A and i^I depend upon the firm's current state. Finally, we turn to the evolution of the firm distribution, $\mu^{'} = \Gamma(\gamma_A, \mu)$. It is useful to define the indicator function $J(x) = 1$ if x is true, and 0 otherwise. For each $(\varepsilon_m, z, k) \in S$

$$\begin{aligned} & \mu^{'}(\varepsilon_m, z, k) \\ = \; & \sum_{l=1}^{N_\varepsilon} \pi^{\varepsilon}_{lm} \left[\begin{array}{c} \int J(\xi \leq \xi^T(\varepsilon, z, k; \gamma_A, \mu)) G\left[\xi^T(\varepsilon_l, z, k; \gamma_A, \mu)\right] \mu(\varepsilon_l, z, dk) \\ + \int \left[1 - \int G\left[\xi^T(\varepsilon_l, z, k; \gamma_A, \mu)\right] \mu(\varepsilon_l, z, dk)\right] J(\xi > \xi^T(\varepsilon, z, k; \gamma_A, \mu)) \end{array} \right] \end{aligned}$$

D Computational Details - Aggregate Uncertainty

D.1 Solution Algorithm

To solve the model, we follow the approach in Khan and Thomas (2008). 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, two endogenous distributions for the capital stocks and the vintage technologies are available to the firms. In theory, this could complicate the solution algorithm by requiring agents to forecast the behavior of two distributions rather than one. In practice, when the persistence of the shock is relatively low, the standard rule that uses the mean of the capital stocks as a regressor works very well, yielding an accurate forecast of prices and future proxy aggregate state. We forecast the mean capital K' and the marginal utility of consumption p using $log(K')$=β_0 + $\beta_1 log(K)$ + ε_K and $log(p)$=β_0 + $\beta_1 log(K)$ + ε_p. As we approximate $\gamma_{A,t}$ using the discretization procedure in Tauchen (1986) using a grid with five points, we estimate a regression conditional on each realization of the aggregate process, $\gamma_{A,t}$. In Tables $A.2$ and $A.3$, we assess the accuracy of the forecasting rule for both models.

Table $A.2$: Forecasting Rules - Vintage Model

	Technology	β_0	β_1	$S.E.$	Adj. R^2
	(A)	(B)	(C)	(D)	(E)
Forecasting K'					
	$\gamma_{A,1}$	0.01613	0.90762	0.00038	0.99948
	$\gamma_{A,2}$	0.00953	0.90435	0.00041	0.99940
	$\gamma_{A,3}$	0.00283	0.90364	0.00041	0.99947
	$\gamma_{A,4}$	-0.00382	0.90251	0.00041	0.99937
	$\gamma_{A,5}$	-0.01050	0.90341	0.00041	0.99941
Forecasting p					
	$\gamma_{A,1}$	0.54530	-0.55812	3.65e-05	0.99998
	$\gamma_{A,2}$	0.54597	-0.55717	3.50e-05	0.99998
	$\gamma_{A,3}$	0.54672	-0.55738	3.62e-05	0.99998
	$\gamma_{A,4}$	0.54745	-0.55771	3.41e-05	0.99998
	$\gamma_{A,5}$	0.54818	-0.55818	3.56e-05	0.99998

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

	Technology	β_0	β_1	$S.E.$	Adj. R^2
	(A)	(B)	(C)	(D)	(E)
Forecasting K'					
	$\gamma_{A,1}$	0.10909	0.76037	2.26e-05	0.99999
	$\gamma_{A,2}$	0.10386	0.76046	1.93e-05	0.99999
	$\gamma_{A,3}$	0.09861	0.76052	2.22e-05	0.99999
	$\gamma_{A,4}$	0.09329	0.76080	1.98e-05	0.99999
	$\gamma_{A,5}$	0.08765	0.76184	2.25e-05	0.99999
Forecasting p					
	$\gamma_{A,1}$	0.60263	-0.30035	8.46e-06	0.99998
	$\gamma_{A,2}$	0.60136	-0.30086	7.45e-06	0.99999
	$\gamma_{A,3}$	0.60013	-0.30148	8.33e-06	0.99999
	$\gamma_{A,4}$	0.59888	-0.30202	7.88e-06	0.99998
	$\gamma_{A,5}$	0.59769	-0.30272	6.63e-06	0.99998

We find that the algorithm yields a very accurate solution as testified by the high R^2 and small standard errors. As discussed by Den Haan (2010), R-squares are averages and scaled by the variance of the dependent variable. To provide a robust statistic, we report the maximum forecast error for each regression. For the vintage model, the maximum percentage errors are 0.016 percent for p and 0.015 for K'. For the RBC model, the maximum percentage errors are 0.006 percent for p and 0.014 percent for for K'. We conclude that the forecasting rules are extremely precise.

E Business Cycle Moments: Robustness

Here, we report the business cycle moments of the economy with vintage technology and its neoclassical counterpart with technology adjustment costs equal to zero.

Table $A.4$: Business Cycle Statistics Technology Shock:
Growth Rates HP-filtered

	ΔGDP	ΔConsumption	ΔInvestment	ΔTFP
	(A)	(B)	(C)	(D)
Data				
σ_X	1.587	1.087	4.431	1.129
$Corr(X, GDP)$	1	0.671	0.435	0.841
Baseline vintage				
σ_X	0.751	0.085	4.883	0.352
$Corr(X, GDP)$	1	0.781	0.991	0.988
Neoclassical model $\overline{\xi} = 0$				
σ_X	0.585	0.127	2.892	0.236
$Corr(X, GDP)$	1	-0.344	0.989	0.919

Note: Each entry represents the volatility of the respective variable. Δ indicates the growth rate. Given the annual frequency of the dayta, we employ the convenitional HP penalty of 100.