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

Giuseppe Fiori Filippo Scoccianti

January 31, 2026

Abstract

We explore how firm heterogeneity affects business cycle dynamics within a general equilibrium model when technological progress is investment-specific. Our data show that investment age—the time elapsed since the firm experienced a significant investment episode—reflects technology obsolescence and is negatively correlated with firms' total factor productivity. In the model, new technologies are embodied in new capital goods, and non-convex adoption costs limit firms' access to new technologies. Idiosyncratic and aggregate shocks alter the technology adoption timing, shifting the cross-sectional distribution of embodied technology and generating endogenous procyclical productivity movements beyond the exogenous shocks to frontier technology. This mechanism, absent in standard neoclassical growth models where all firms instantaneously adopt the latest technology, amplifies aggregate dynamics compared to a neoclassical growth model.

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

*Giuseppe Fiori: Board of Governors of the Federal Reserve System, Division of International Finance, 20th and C St. NW, Washington D.C. 20551, United States (email: giuseppe.fiori@frb.gov; webpage: http://www.giuseppefiori.net); Filippo Scoccianti: Banca d'Italia, Via Nazionale 91, Rome, Italy (email: filippo.scoccianti@bancaditalia.it). We thank the Editor, Vincenzo Quadrini, and three anonymous referees for helpful comments that improved the quality of the manuscript. We are also grateful to Andrea Caggese, Federico Cingano, Marcela Eslava, Jesús Fernández-Villaverde, Domenico Ferraro, Reint Gropp, Michael Koetter, Massimo Giovannini, David Jenkins, 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

Standard models of firm heterogeneity treat productivity differences as purely random shocks (Hopenhayn, 1992), while models of investment-specific technological change assume all firms instantaneously access the technological frontier (Greenwood, Hercowitz and Krusell, 1997). Yet neither framework captures a fundamental feature of technology diffusion: firms adopt new technologies at different times through lumpy capital investments, creating systematic productivity differences based on the timing of adoption rather than pure randomness or instantaneous access. This pattern is particularly salient in recent waves of technological change, such as the adoption of artificial intelligence and automation technologies, where substantial heterogeneity exists in both the timing and extent of firm-level implementation (McElheran et al., 2024).

This paper bridges these literatures by examining how capital-embodied technological progress generates productivity heterogeneity and shapes business cycle dynamics. We establish two key empirical findings. First, we show that investment age proxies for technology vintage: firms report operating more obsolete technology when investment age is higher, and measured TFP declines by approximately 0.5% per year of investment age. This relationship survives extensive controls including lagged TFP, firm fixed effects, time-varying sectoral trends, and managers' forward-looking sales expectations. Second, this relationship operates causally: sectors experiencing identified investment-specific productivity shocks exhibit increased investment rates, reduced investment age, and rising TFP with a one-period lag.

We then develop a general equilibrium model building on Khan and Thomas (2003, 2008) where non-convex adoption costs prevent firms from continuously accessing the technological frontier. As firms invest in the latest technology at different times, the economy features a non-degenerate distribution of capital vintages. Crucially, aggregate and idiosyncratic shocks alter technology adoption timing, shifting the cross-sectional distribution of embodied technology. This generates endogenous procyclical productivity movements beyond the exogenous shocks to frontier technology—a mechanism absent in standard neoclassical growth models where all firms instantaneously adopt the latest technology.

Our quantitative analysis reveals that microeconomic heterogeneity in vintage technol-

ogy substantially amplifies business cycles. Following a one-standard-deviation investment-specific productivity shock, the model with vintage heterogeneity amplifies the volatility of output by approximately 50%. This finding contrasts sharply with Khan and Thomas (2003, 2008), where firm heterogeneity proves irrelevant to aggregate dynamics due to offsetting general equilibrium effects. In their framework, all firms access identical technology, so heterogeneity reflects only capital quantity. In our framework, firms choose when to upgrade from their current vintage to more efficient investment goods. This mechanism generates endogenous heterogeneity in technology quality: as laggard firms operate increasingly obsolete vintage—falling behind the frontier at a positive rate for each period of non adoption—productivity gaps widen, strengthening upgrade incentives. When an investment-specific productivity shock improves the frontier technology, more firms find it profitable to upgrade, shifting the composition of the vintage distribution toward more efficient technologies. This compositional effect generates persistent productivity improvements beyond the direct impact of the shock, amplifying aggregate dynamics and delivering greater persistence of macroeconomic series. Our contribution differs fundamentally also from Winberry (2021) and Koby and Wolf (2019): while these papers focus on heterogeneity in capital quantity arising from adjustment costs, we introduce heterogeneity in capital quality—the efficiency of investment goods themselves—which generates amplification through compositional shifts in the vintage distribution that create endogenous productivity movements, even when operating with the elasticity of investment demand with respect to the real interest rate being similar to Khan and Thomas (2008).

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

Using data from the Survey of Industrial and Service Firms (INVIND), we show that

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

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 after controlling for time fixed effects, industry-time fixed effects, and firm fixed effects, as well as firm age, size, lagged TFP, and managers' expectations about future sales. This microeconomic evidence supports investment-specific technological progress and is inconsistent with models where TFP is purely exogenous, as in Hopenhayn (1992) or Khan and Thomas (2008). The vast heterogeneity in investment age in the data points to widespread differences in technologies available to firms, with the average firm displaying an investment age equal to about three and a half years.

To study the role of investment age heterogeneity for aggregate dynamics, we formulate a general equilibrium framework that builds on Khan and Thomas (2003, 2008) where new technologies are embodied in new investment goods and firms are subjected to exogenous idiosyncratic shocks. The distinctive feature of the model is a quality ladder structure in investment-specific productivity. In the model, the firm's productivity, which maps physical units of investment on to capital services, is endogenous to the timing of technology adoption. Every period, firms optimally decide when to invest in the latest vintage of capital goods and, therefore, the associated latest technology. As this choice is subject to a non-convex adoption cost, the firm's policy functions are of the (S,s) type: Some firms buy the newest capital goods and adopt the latest technology, while others postpone it and keep buying the current vintage they operate. Conditional on the adoption decision, firms optimally choose investment. In equilibrium, technologies of different quality coexist. The history of adoption choices contributes to microeconomic heterogeneity in productivity and, in turn, determines aggregate productivity. Instead, when the adoption cost is set equal to zero, all the firms find it optimal to adopt the latest technology in every period, with our framework boiling down to a standard neoclassical growth model with investment-specific technical progress as in Greenwood, Hercowitz and Krusell (1997, 2000), where firms' and aggregate productivity are exogenous. What differentiates our model from Khan and Thomas (2003), Solow (1960) and Greenwood, Hercowitz and Krusell (2000), is the ability of firms to invest in capital goods embodying different efficiencies. Whereas prior work examines variation in capital quantity induced by adjustment frictions while abstracting from techno-

logical heterogeneity, we analyze variation in capital quality–the productivity of investment goods–which amplifies shocks through a novel channel: changes in the vintage composition generate endogenous fluctuations in aggregate productivity that magnify business cycle dynamics beyond pure accumulation effects.

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

The main result of the paper is that microeconomic heterogeneity is a key determinant of business cycle dynamics following investment-specific productivity shocks. In response to an exogenous increase in the efficiency of the frontier, the canonical investment-specific shock, firms increase investment by acquiring capital embodying a more efficient technology. The ensuing increase in aggregate productivity is larger than the one implied by the shock as firms switch from less efficient capital goods to more efficient ones. From a quantitative standpoint, this channel boosts aggregate productivity, amplifying aggregate series' persistence and volatility relative to a standard neoclassical growth model with investment-specific technological progress. Our findings contrast with the earlier work by Thomas (2002) and Khan and Thomas (2003, 2008), where including firm heterogeneity is irrelevant for aggregate dynamics relative to a neoclassical benchmark because of general equilibrium effects that dampened the increase in investment. In our framework, general equilibrium effects do not neutralize firm heterogeneity because of the productivity gains associated with the investment decision. As the latest technology is more efficient than the previous one, the current technology available to the firm becomes more obsolete relative to the technological frontier, increasing the benefit of adoption. As more firms adopt the latest vintage, the share of firms operating higher-productivity technology remains above its steady state leading to a persistent increase in the economy-wide productivity. The endogenous response of productivity constitutes an additional force that amplifies fluctuations in investment and macroeconomic dynamics relative to a neoclassical growth model. For a given persistence, the vintage technology model requires shocks about 50% smaller than the neoclassical growth

model to generate business cycles of the same magnitude. Our results support the view that microeconomic heterogeneity is relevant for aggregate dynamics as investment contributes to productivity.

The paper proceeds in three parts. Section 1.1 describes our contribution relative to the existing literature on lumpy investment and investment-specific technological progress. Sections 2–3 establish the empirical relationship between investment timing and productivity. Section 2 describes data covering the near-universe of Italian firms and constructs investment age as a measure of technology vintage. Section 3 documents that investment age negatively predicts TFP in firm-level panel regressions with extensive controls and demonstrates causality using sectoral responses to SVAR-identified investment-specific shocks. Sections 4–6 develop and analyze the model. Section 4 formulates a general equilibrium framework with optimal technology adoption. Section 5 calibrates the model, and Section 6 validates its fit to microeconomic distributions of investment rates, investment age, and TFP dispersion.

Section 7 quantifies the amplification mechanism through impulse response analysis and business cycle accounting, showing that vintage heterogeneity substantially strengthens the propagation of investment-specific technology shocks. Section 8 concludes.

1.1 Literature Review

This paper bridges three literatures that have largely developed independently: empirical work on capital-embodied technological change, theoretical models of lumpy investment and firm heterogeneity, and business cycle dynamics with investment-specific technology shocks. We show that integrating insights from these literatures generates new predictions about both microeconomic behavior and aggregate dynamics.

Empirical evidence on capital-embodied technology at the firm level. While capital-embodied technological progress is central to growth theory (Solow, 1960; Greenwood, Hercowitz and Krusell, 1997), systematic firm-level evidence remains scarce. Most studies rely on aggregate or sectoral data (Gordon, 1990, Hulten, 1992; Wolff, 1996; Cummins and Violante, 2002), using investment price deflators as indirect measures of embodied technical change.[2] The few firm-level studies reach conflicting conclusions: Power (1998) finds no ev-

[2] Knittel (2011) and Bertolotti, Gavazza and Lanteri (2023) estimate the evolution of the technological frontier

idence that investment spikes increase productivity in U.S. manufacturing, while Sakellaris and Wilson (2004) and Licandro, Maroto Illera and Puch (2005) find the opposite using U.S. and Spanish data respectively. We advance this literature in three ways. First, we construct a direct firm-level proxy for technology vintage-investment age-that does not rely on price deflators and that firms themselves validate through self-reported technology assessments. Second, we address identification concerns that have plagued previous work by controlling for forward-looking expectations (using survey-elicited sales forecasts) and by exploiting SVAR-identified aggregate shocks. Third, we connect firm-level evidence to aggregate implications through a calibrated general equilibrium model, moving beyond reduced-form documentation to structural interpretation.

Lumpy investment and aggregate dynamics. An extensive literature examines whether microeconomic non-convexities in capital adjustment aggregate to affect business cycles. Caballero and Engel (1999) and Bachmann, Caballero and Engel (2013) argue that adjustment cost heterogeneity amplifies aggregate shocks, while Thomas (2002) and Khan and Thomas (2008) demonstrate that general equilibrium effects can neutralize this amplification, rendering firm heterogeneity irrelevant for aggregate dynamics—the "irrelevance result."[3] More recent work by Winberry (2021) and Koby and Wolf (2019) reconcile these findings by showing that amplification depends critically on the elasticity of investment demand with respect to the real interest rate: Khan–Thomas calibrations imply implausibly high elasticities, while more realistic values generate amplification in investment relative to a neoclassical growth benchmark.

Our contribution differs fundamentally from this debate. These papers focus on heterogeneity in capital *quantity* arising from non-convex adjustment costs, abstracting from technology differences. We introduce heterogeneity in capital *quality*—the efficiency of investment goods themselves—which generates amplification through a distinct mechanism:

in cars without using prices and by positing a marginal-cost function that depends on vehicle attributes.

[3]Other important contributions include Cooper and Haltiwanger (1993), Fiori (2012), and House (2014), just to name a few. Cooper and Haltiwanger (1993) study the aggregate Implications of machine replacement with an application to the auto industry. Fiori (2012) shows that a two-sector model calibrated to plant-level investment patterns naturally generates the hump-shaped response of aggregate investment to shocks and the positive-then-negative autocorrelation of investment growth observed in U.S. data. House (2014) demonstrates that fixed costs generate realistic aggregate investment dynamics when calibrated to plant-level evidence, though this amplification disappears when investments are sufficiently long-lived, as the cross-sectional distribution then becomes irrelevant for equilibrium dynamics.

compositional shifts in the vintage distribution create endogenous productivity movements. This mechanism operates even with investment demand elasticities consistent with Khan and Thomas (2008), because productivity gains from upgrading to more efficient technologies amplify shocks beyond what occurs through capital accumulation alone.

Business cycle implications of investment-specific technology. In estimated DSGE models, investment-specific technology shocks explain 30–60% of output volatility (Smets and Wouters, 2007; Justiniano, Primiceri and Tambalotti, 2010). Yet standard implementations assume all firms instantaneously adopt the technological frontier, abstracting from diffusion dynamics. We show that accounting for heterogeneous adoption timing substantially strengthens the propagation of these shocks: the vintage model requires investment-specific innovations approximately 50% smaller than the standard model to match observed volatility, implying previous estimates may overstate the true variance of technology shocks.

Our framework also generates countercyclical TFP dispersion endogenously—more firms operate near the frontier during booms—without relying on time-varying volatility of firm-level shocks (Bloom et al., 2018; Bachmann and Bayer, 2014). This suggests an alternative interpretation of the "uncertainty shocks" literature: observed increases in productivity dispersion during recessions may partly reflect compositional effects from delayed technology adoption rather than purely exogenous increases in the variance of firm-level shocks.

Vintage capital models. Our work revives the vintage capital literature initiated by Solow (1960) and Johansen (1959), which emphasized that technological progress is embodied in new capital goods. While influential in growth theory, this literature faced important theoretical challenges due to aggregation difficulties: without knife-edge conditions (e.g., constant returns to scale), the economy cannot be represented by a single aggregate capital stock, complicating both theoretical analysis and empirical implementation (for a survey on the seminal contributions in this literature, see Boucekkine, de la Croix and Licandro, 2011).

We overcome these difficulties through computational methods developed in the heterogeneous agents literature. By explicitly tracking the distribution of firms over capital stocks and vintages, we can study vintage capital in a general equilibrium framework with realistic micro-level heterogeneity. This allows us to confront the model with both microeconomic data (investment rate distributions, TFP dispersion) and aggregate moments (business cycle statistics), providing discipline absent from earlier vintage capital work.

Intangible capital and technology embodiment. While our primary focus is physical capital (equipment and structures), our empirical results (Table 1, Columns C–D) reveal that investment age for intangible capital (R&D, software, organizational capital) also negatively predicts TFP. This suggests embodiment is a general feature of capital investment rather than specific to physical capital. Recent work emphasizes the role of R&D expenditures (Wang and Zhang, 2025) and complementarities between physical capital, intangibles, and skilled labor in generating productivity gains (Bresnahan, Brynjolfsson and Hitt, 2002; Bloom, Sadun and Van Reenen, 2012). Incorporating these complementarities could amplify our mechanism, as firms adopting new physical capital optimally invest in complementary intangibles and reorganize production. We leave this extension to future work, but note that our estimates likely provide a lower bound on the importance of embodied technological change.

2 Firm Heterogeneity in Investment Age and Technology

Technological progress is embodied in capital goods, making the vintage of capital embodied in firms' capital stocks a key determinant of cross-firm productivity differences; see, for instance, Solow (1960) and Greenwood, Hercowitz and Krusell (1997). Because firms acquire new technology through discrete investment episodes rather than continuous upgrades, the timing of these investments provides a natural measure of technological vintage. This section exploits this feature—the lumpy, infrequent nature of firm-level investment—to construct a proxy for the technology available to firms. We define *investment age* as the time elapsed since a firm's last substantial capital investment episode and demonstrate that it captures meaningful variation in firms' proximity to the technological frontier. After describing our data sources in Section 2.1 and our measurement of capital and productivity in Section 2.2, we validate investment age as a technology proxy using survey evidence in Section 2.3. Section 2.4 then documents substantial cross-firm heterogeneity in investment age, revealing widespread differences in capital vintage. This heterogeneity provides the empirical foundation for our subsequent analysis of how the timing of capital accumulation shapes firm-level productivity through technology adoption and diffusion.

2.1 Description of Data Sources

We combine three primary data sources to construct our firm-level measures of investment and productivity. The core variables—payroll, gross value-added, and employment—come from yearly balance sheets in the Cerved Group S.P.A. database (Cerved), supplemented by the Bank of Italy's annual INVIND survey and the Italian National Institute for Social Security (INPS) (see Appendixes A and B for detailed information on data sources and variable construction).

Cerved provides comprehensive balance sheet information for Italian firms over the period 1986–2016. The database covers up to 80% of Italian value added and closely matches the size and sectoral distribution of the Italian firm population.[4]

We supplement Cerved with INVIND, an annual business survey conducted by the Bank of Italy that collects information on firms' expectations, investment decisions, and available technology. The survey targets a representative sample of firms operating in industrial sectors (manufacturing, energy, and extractive industries), construction, and nonfinancial private services, covering 1994–2016.[5] INVIND serves two critical roles in our analysis. First, it contains firms' self-reported assessments of their operating technology, which we use to validate investment age as a technology vintage proxy. Second, it includes managers' expectations about future sales, allowing us to control for forward-looking factors that might confound the relationship between investment age and TFP.

2.2 Measurement of TFP and Capital

The adoption of breakthrough technologies—from computerization to artificial intelligence—requires firms to invest substantially in capital goods that embody these advances (Solow, 1960; Johansen, 1959; Greenwood, Hercowitz and Krusell, 1997). Linking investment to productivity dynamics is therefore central to understanding how new technologies diffuse across firms and shape aggregate productivity growth.

[4]Table A.1 reports the sectoral composition of our sample, classified according to NACE (the statistical classification of economic activities in the European Community). Consistent with their importance in the Italian economy, manufacturing and trade sectors account for more than half of all observations.

[5]Specifically, INVIND ensures representativeness across branch of activity (according to an 11-sector classification), size class, and region of the firm's headquarters.

Linking investment to productivity requires measuring how capital vintage affects productive efficiency. Following the empirical literature, we construct real capital stocks using the perpetual inventory method with sector-specific depreciation rates from ISTAT. This standard approach measures capital in consumption units—the cost of acquiring capital goods—rather than efficiency units adjusted for vintage quality. Consequently, our TFP measure combines neutral productivity (affecting all inputs equally) with investment-specific productivity embodied in capital goods, as shown by Greenwood, Hercowitz and Krusell (1997, Appendix B). We emphasize that this is a feature, not a limitation. When capital is measured in consumption units, productivity improvements from adopting more efficient capital goods appear as increases in measured TFP. This makes investment age-the time since acquiring new capital-a natural proxy for embodied technology: firms with higher investment age have older, less efficient capital embodied in their capital stock, which manifests as lower measured TFP. Section 3 validates this interpretation empirically.

This measurement constraint leads us to adopt the perpetual inventory method—the standard approach in the empirical literature—to construct real capital stocks. We apply this method to balance sheet data from our sample firms, aggregating all capital goods using physical depreciation rates without vintage adjustments (see Appendix B for details). We deflate investment in tangible assets using sectoral deflators for capital expenditures and employ sector-specific depreciation rates from the Italian National Institute of Statistics (ISTAT). Because balance sheet data do not separately report equipment and structures, our investment measure combines both categories—a limitation that prevents us from exploiting their different depreciation patterns and technology content. Separate treatment could sharpen identification, as investment age for equipment should predict TFP more strongly than investment age for structures. While some studies use investment composition as an instrument (Sakellaris and Wilson, 2004), limited data prevent this approach. To the extent equipment investment drives our results, our estimates may understate the importance of embodied technology, as aggregation with slower-depreciating structures attenuates the measured effect.

2.3 Investment Age As a Proxy for Firm's Technology

To proxy for the vintage of technology embodied in firms' capital stocks, we exploit a fundamental characteristic of firm-level investment: its lumpiness. Investment at the firm level does not occur smoothly but rather occurs in large, infrequent episodes followed by periods of subdued activity. Following the existing literature and recently Bachmann and Bayer (2014), we classify firms as "lumpy adjusters" when they exhibit a positive investment spike—defined as an investment rate exceeding 20% of the capital stock. We adopt this threshold following convention in the lumpiness literature, which has shown that this threshold effectively identifies economically meaningful capital expenditures (Gourio and Kashyap, 2007; Bachmann and Bayer, 2014) providing a time-varying measure of how long each firm has operated without major capital renewal.

Lumpy investment is pervasive in our data, consistent with extensive evidence from other economies. As reported in Appendix C, in any given year, approximately 20% of firms in our sample experience investment spikes, yet these spike investors account for roughly two-thirds of aggregate investment. This pattern aligns with documented findings from the United States (Doms and Dunne, 1998), Germany (Bachmann and Bayer, 2014), Spain (Licandro, Maroto Illera and Puch, 2005), Norway (Øivind and Schiantarelli, 2003), and Chile (Gourio and Kashyap, 2007).

We construct our firm-level measure of investment age, $\text{Inv.Age}_{f,t}$, based on the timing of these investment spikes. Specifically, $\text{Inv.Age}_{f,t}$ measures the years elapsed since a firm's most recent *positive* investment spike. At the moment of a spike, $\text{Inv.Age}_{f,t}$ resets to zero and subsequently increments by one each year until the firm experiences another positive spike. This measure serves as a proxy for the technological vintage embodied in the firm's capital stock, operationalizing the notion that technological progress is embodied in new capital goods.

Constructing this measure requires observing a sufficiently long sample period to initialize the distribution and apply panel estimation techniques. We therefore allocate approximately 40% of the sample period to initialize the investment age distribution (1986–1995 for Cerved data), reserving 1996–2016 for our main empirical analysis. Note that because the Cerved and INVIND datasets span different time periods, the specific sample windows

differ across data sources.

Investment Age and Vintage Technology We employ INVIND survey data to validate the interpretation of $Inv.Age$ as a measure of vintage technology—specifically, firms' distance from the technological frontier. In its 2014 wave, INVIND asked firms to self-assess the advancement of their production technology using a discrete index from 4 (most advanced technology) to 1 (obsolete technology).

Using the 2014 cross-section, we find a negative relationship: Firms with higher investment age employ less advanced technologies. Specifically, $Tech.Adv_{f,2014} = 3.529 - 0.014 \times Inv.Age_{f,2014} + \epsilon_{f,2014}$, where the coefficient on investment age is statistically significant at 5% (p-value equal to 0.025) and the constant is significant at 1%.[6] This cross-sectional evidence supports the interpretation of $Inv.Age_{f,t}$ as a firm-level measure of the distance from the technological frontier.

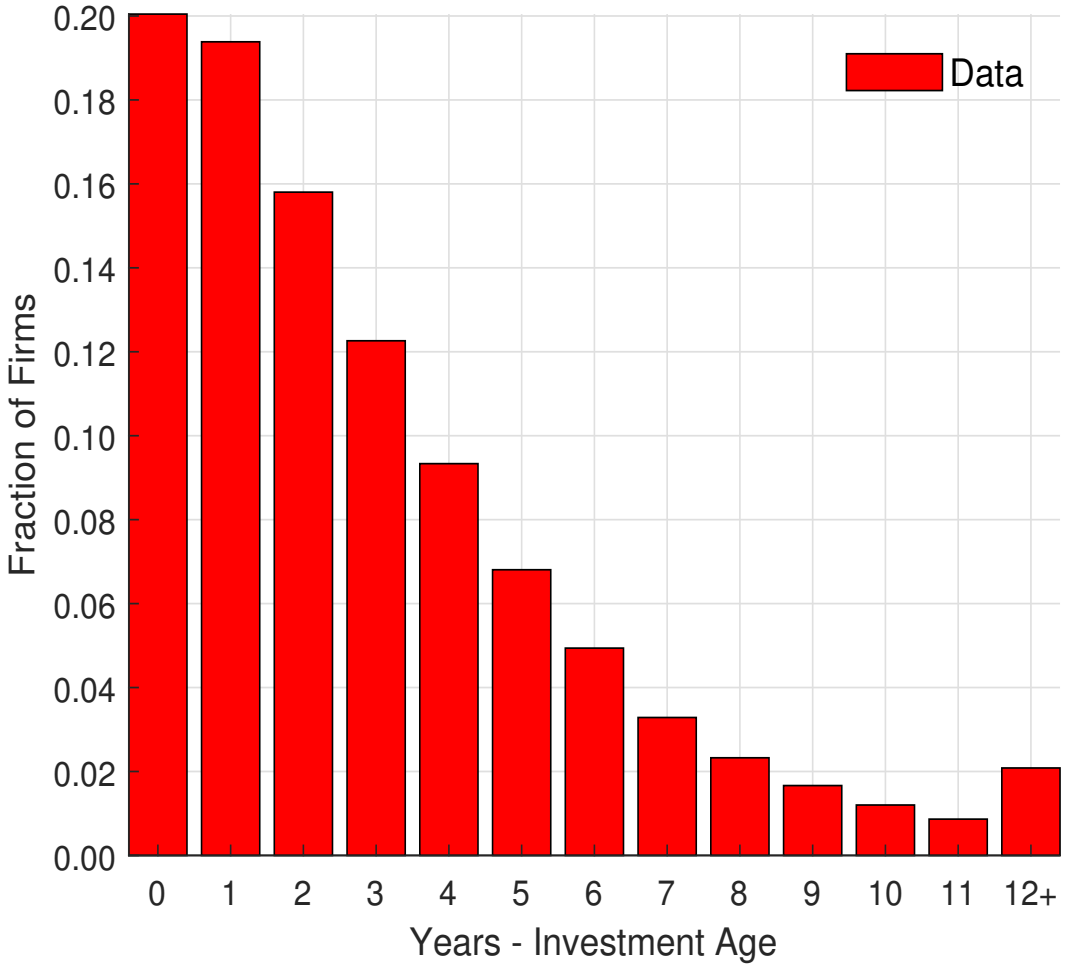
2.4 The Cross-Sectional Distribution of Investment Age

The cross-sectional distribution of investment age in the data, obtained by averaging across the sample period, points to substantial heterogeneity in firms' investment age and indicates widespread differences in technologies available to firms (see Figure 1). As described in the previous section, the fraction of firms that display an investment age equal to zero has experienced an investment spike in the current period. As firms delay a significant capital adjustment, investment age progressively increases. Average investment age is 3.5 years, indicating the typical firm undertakes major capital expenditures every 3–4 years. Dispersion is substantial: approximately 20% of firms spike in any given year ($Inv.Age_{f,t} = 0$), while 15% have investment age exceeding 6 years, revealing considerable vintage heterogeneity across the economy.

Investment age is strongly procyclical. During the early 2000s expansion, approximately 22% of firms spiked annually, reducing average investment age to 3.2 years. During the 2009 Global Financial Crisis and 2012 Sovereign Debt Crisis, the spike rate fell to 11%. This cyclical pattern—firms delaying major capital upgrades during downturns and accelerating them during booms—suggests that technology adoption is endogenous to economic condi-

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

Figure 1: Investment Age Distribution



tions, a mechanism central to our model in Section 4. 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

Having established investment age as a valid proxy for technology vintage and documented substantial heterogeneity in this measure across firms, we now examine its implications for productivity. Our central empirical finding is that firms with higher investment age—those further from their last major capital investment—exhibit systematically lower TFP. This negative relationship provides direct evidence that technological progress is embodied in capital goods: as firms delay investment, they fall further behind the technological frontier,

depressing measured productivity.

We establish this result in two stages. Section 3.1 documents a robust negative correlation between the lag of investment age and TFP in firm-level panel regressions. Section 3.2 strengthens the causal interpretation by exploiting aggregate investment-specific shocks identified following Fisher (2006). Sectors experiencing positive shocks exhibit systematic increases in investment rates followed by TFP gains, confirming that technology embodied in new capital causally affects productivity. These findings validate the central mechanism of capital—embodied technological progress underlying our theoretical framework in Section 4.

3.1 Micro Evidence: Investment Age and TFP

Our empirical strategy exploits within-firm variation in investment timing to identify the relationship between technology vintage and productivity. We estimate the following baseline specification:

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

where $TFP_{f,t}$ is firm f's total factor productivity in year t, measured as the Solow residual from a Cobb-Douglas production function. Following Bachmann and Bayer (2014), we estimate output elasticities as the median factor expenditure shares within each industry (see Appendix B for details). As established in Section 2.2, this TFP measure captures both neutral and investment-specific technological change. The coefficient β captures the correlation between lagged investment age and current productivity. We lag investment age by one period to account for installation and adjustment costs that delay the productivity impact of new capital—a specification consistent with standard time-to-build assumptions in the investment literature

The specification includes a comprehensive set of controls: lagged TFP (to capture productivity persistence), firm age and size (lifecycle and scale effects), firm fixed effects (time-invariant heterogeneity), industry fixed effects, time fixed effects (aggregate shocks), and industry-time fixed effects (sector-specific trends). This demanding specification identifies β

from within-firm variation in investment age around firm-specific means, while simultaneously netting out common aggregate shocks and differential sectoral trends.

We emphasize that our estimates reflect predictive relationships rather than causal effects. Time-varying firm-specific factors—such as anticipated profitability shocks or persistent productivity trends—could simultaneously influence both current TFP and investment decisions, generating spurious correlations. The INVIND sample partially addresses these concerns by including managers' forward-looking sales expectations ($s^e_{avg,f,t}$), allowing us to control for information that might otherwise confound the investment age-TFP relationship.

Table 1: Investment Age and Total Factor Productivity

	$\frac{TFP_{f,t}}{}$ (A)	$\frac{TFP_{f,t}}{}$ (B)	$\frac{TFP_{f,t}}{}$ (C)	$\frac{TFP_{f,t}}{}$ (D)	$\frac{TFP_{f,t}}{}$ (E)	$\frac{TFP_{f,t}}{}$ (F)
$Inv.Age_{f,t-1}$	-0.439^{***}	-0.436^{***}	-0.335^{***}	-0.335^{***}	-0.759^{*}	-0.844^{**}
	(0.01)	(0.01)	(0.01)	(0.01)	(0.395)	(0.414)
$Inv.Age^{Int.}_{f,t-1}$			-0.157^{***}	-0.157^{***}		
			(0.01)	(0.01)		
$s^e_{avg,f,t}$						0.305^{***}
						(0.070)
N. of obs.	1,837,323	1,795,163	1,345,848	1,341,856	1,487	1,234
R^2	0.710	0.712	0.717	0.718	0.936	0.933
Estimator	OLS	OLS	OLS	OLS	OLS	OLS
Firm FE	✓	✓	✓	✓	✓	✓
Industry FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Time×Ind. FE	✓	✓	✓	✓		
Source of Data	Cerved				INVIND	
Sample Period	1998-2016	1998-2016	1998-2016	1998-2016	2003-2015	2003-2015

Note: * p<0.10, ** p<0.05, and *** p<0.01. Standard errors are 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 in physical capital, defined as the firm experiencing an investment rate above 20% at time $t-1$. $Inv.Age^{Imm}_{f,t-1}$ measures the time elapsed between investment spikes in intangible capital, defined as the firm experiencing an investment rate above 60% at time $t-1$. $s^e_{avg,f,t}$ denotes managers' subjective expectations about sales one year ahead. Columns A through D report estimates obtained using Cerved data. Columns E and F report estimates obtained using INVIND data.

Table 1 documents a robust negative relationship between investment age and TFP across all specifications. Column A establishes our baseline result using the universe of firms in Cerved (1.8 million firm-year observations spanning 1998-2016). The coefficient on investment age is negative and significant at 1% level, with an estimated productivity loss of 0.44% per year. This supports our hypothesis that firms operating older vintages of capital fall behind the technological frontier. Column B addresses concerns that entry and exit dynamics might drive this relationship. Restricting the sample to firms at least three years old yields virtually identical coefficients (-0.436 versus -0.439), confirming that the investment age-TFP relationship reflects within-firm dynamics rather than selection effects.

Columns C and D impose a more stringent test by controlling for intangible capital investment timing through $Inv.Age^{Int}_{f,t-1}$.[7] This addresses the concern that our results might simply reflect correlated investment patterns across different asset types rather than embodied technological change in physical capital specifically. While the negative coefficient on intangible investment is significant, physical capital investment age somewhat smaller (-0.335 versus -0.439) remains economically meaningful and statistically significant at the 1% level. Importantly, this demonstrates that even after accounting for firms' investment timing in intangible assets, physical capital vintage retains substantial independent explanatory power for TFP dynamics. This result is consistent with capital-embodied technological progress operating operating through physical capital specifically, as our model in Section 4 emphasizes.

Columns E and F turn to the INVIND sample, which enables the most demanding robustness check: controlling for managers' forward-looking expectations about future sales. This is crucial because expectations about future business conditions could simultaneously drive both current investment decisions and future productivity outcomes, potentially generating a spurious correlation between investment age and TFP. Although the sample shrinks sig-

[7] We construct the stock of intangible capital (comprised mostly of R&D and software expenditures) by applying the perpetual inventory method to real intangible investment following the same approach used for physical investment goods (equipment and structures). We assume a depreciation rate of 20%, which is consistent with the existing literature. Despite substantial heterogeneity across asset types and industries, several benchmark rates have emerged as standard practice. For R&D capital, the most commonly used depreciation rate is approximately 15%, though estimates range from 12% (Nadiri and Prucha, 1996) to 20% (Lev and Sougiannis, 1996) depending on methodology and sample period. In the absence of guidance on what constitutes a spike in intangible investment, we experiment with a threshold of 60% as in Wang and Zhang (2025) to deliver similar investment age dynamics as for physical capital. Alternative spike definitions between 20% and 80% do not alter our conclusions.

nificantly, Column E presents results for the INVIND sample without expectations, showing a coefficient of -0.759 (significant at 10%). Column F adds expected sales growth, which enters positively and highly significantly, confirming that forward-looking information indeed predicts TFP. Critically, even with this forward-looking information controlled for, investment age maintains a negative and statistically significant coefficient of -0.844, the coefficient is larger but less precisely estimated (now significant at 5%). Despite controlling for information that managers themselves use to forecast future performance, our results show that capital vintage effects operate independently of expectations-driven productivity dynamics.

Several robust patterns emerge across specifications. First, the negative investment age-TFP relationship holds consistently across two independent datasets (Cerved and INVIND), different sample restrictions, and progressively more demanding control sets. Second, the estimated coefficients remain economically meaningful, averaging approximately -0.4 to -0.8% per year of investment age—effects that accumulate to substantial productivity gaps over the multi-year investment cycles documented in Section 2.4. Third, the relationship survives controls for the most plausible confounding factors: intangible investment patterns, firm-specific trends, aggregate and sectoral shocks, and forward-looking expectations.

Fourth, we validate our interpretation using simulated data from the structural model in Section 6.4. Estimating equation (1) on data generated by models with investment-specific technological change replicates the negative investment age-TFP relationship. Models driven exclusively by neutral productivity shocks—whether firm-specific or aggregate—fail to generate this pattern because TFP evolves independently of investment timing. Once we condition on lagged TFP, lagged investment age contains no additional information about current TFP in these neutral-shock models. That investment age retains strong predictive power in the data—even controlling for lagged TFP—therefore constitutes direct evidence of embodied technological progress.

Our findings align with broader cross-country evidence: Fiori, Lipa and Wu (2025) document similar negative relationships between investment age and TFP across US and European firms, with coefficients comparable to ours, indicating that embodied technological progress is a fundamental determinant of productivity dynamics in developed economies. We now complement this firm-level evidence with sectoral analysis exploiting exogenous

investment-specific productivity shocks.

3.2 Sectoral Evidence: Investment Age and TFP

To provide a causal interpretation of the firm-level investment age-TFP relationship, we now examine sectoral responses to plausibly exogenous investment-specific productivity shocks. Our identification strategy follows, in spirit, the hybrid approach of Cloyne et al. (2023) and Sedláček (2020), combining aggregate shock identification with micro-level analysis. We first identify investment-specific productivity shocks at the aggregate level using a structural vector autoregression (SVAR) with long-run restrictions, then trace their effects through sectoral investment and productivity dynamics.

Following Fisher (2006) we impose two key identifying assumptions. First, the relative price of investment has a single permanent source of variation: shocks to investment-specific productivity. Second, labor productivity has two permanent sources: shocks to investment-specific productivity and shocks to neutral technology. Transitory shocks to demand, supply, or other factors remain unidentified, as they do not affect long-run levels. Appendix D provides full details on the estimation.

The key advantage of long-run identification is that it does not require specifying contemporaneous relationships between variables–a particularly appealing feature given uncertainty about short-run adjustment dynamics in investment and productivity. The main limitation is relatively wide confidence intervals, typical for SVARs with long-run restrictions. Despite this, we obtain statistically significant responses for key variables, providing confidence in the identification.

This approach allows us to isolate exogenous variation in the pace of embodied technical change and examine whether sectors experiencing positive investment-specific shocks exhibit the predicted pattern: increased investment rates, reduced investment age, and rising TFP. Finding this pattern would provide causal evidence that technology embodied in new capital drives productivity dynamics, validating our interpretation of the firm-level correlations documented in Section 3.1.

Using these identified shocks, we estimate contemporaneous and one year ahead sectoral responses of investment rates, investment age, and TFP via local projections. We adopt

this two-step approach because the limited availability of annual data prevents us from directly including sectoral variables in the VAR without rapidly exhausting degrees of freedom. We consider specifications ranging from a parsimonious model that includes only the investment-specific productivity shock to a more comprehensive specification that incorporates additional controls: contemporaneous log GDP and three lags (to control for aggregate business cycle conditions and given that we cannot include year fixed effects), two lags of the growth rate of the relative investment price (capturing persistence in investment costs beyond the identified shock), and two lags of sectoral TFP (absorbing sector-specific productivity trends).

Table 2: Sectoral Evidence: Investment Age and TFP

	Panel A: Sectoral Investment Age			
	$Inv.Age_t$	$Inv.Age_{t+1}$	$Inv.Age_t$	$Inv.Age_{t+1}$
	(1)	(2)	(3)	(4)
Inv.-specific productivity shock	-0.058^{***}	-0.076^{***}	-0.037^{***}	-0.054^{***}
	(0.01)	(0.01)	(0.01)	(0.01)
R^2	0.649	0.687	0.689	0.724
Observations	949	949	949	949
	Panel B: The Log of Sectoral TFP			
	TFP_t	TFP_{t+1}	TFP_t	TFP_{t+1}
	(1)	(2)	(3)	(4)
Inv.-specific productivity shock	0.072	1.334^{***}	-0.475	0.845^{*}
	(0.33)	(0.34)	(0.38)	(0.44)
R^2	0.966	0.962	0.987	0.978
Observations	949	949	949	949

Notes: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Every equation includes sector-specific fixed effects, the contemporaneous value and first three lags of log GDP, and the first and second lags of the growth rate of the relative price of investment. Columns 3 and 4 also include time-varying sectoral controls: the first and second lags of sectoral TFP.

Table 2 presents the estimated sectoral responses to investment-specific productivity shocks. Panel A reports the effects on investment age. A positive shock—equivalent to a 0.7 percentage point decline in the relative price of investment (one standard deviation)—increases the share of firms experiencing spikes by 2 percentage points on average on im-

pact and one year ahead. This surge in investment activity reduces cumulatively investment age by approximately 0.1 years, depending on the specification, with effects persisting for two years. The mechanism is straightforward: productivity improvements embodied in new capital goods stimulates investment demand. As more firms undertake major capital projects in response to the shock, average sectoral investment age declines, moving the sector closer to the technological frontier.

The sectoral TFP response exhibits a distinct lag relative to investment dynamics. Panel B (columns 1 and 3) shows that the investment-specific shock has no statistically significant impact on sectoral TFP, despite the immediate surge in investment activity. This delayed response is consistent with the time required for new capital to be installed and incorporated into the production process before productivity gains materialize. Once this integration is complete, however, TFP at the sectoral level rises substantially—by approximately 1% one period after the shock—representing economically significant productivity improvements. Investment age continues to decline one period after the shock, reflecting the sustained impact of the initial productivity improvement on capital investment decisions. If the shock reflected neutral productivity improvements, TFP would respond contemporaneously rather than lagging investment.

These sectoral results strengthen our empirical case in two ways. First, they support that the negative investment age-TFP correlation documented at the firm level reflects causal effects of investment-specific productivity shocks rather than spurious correlations. The temporal pattern-investment age declines immediately, TFP rises with a lag-rules out alternative explanations based on neutral productivity shocks or demand-driven investment cycles. Second, the results provide quantitative discipline for structural models: a one-standard-deviation investment-specific shock (0.7 percentage point decline in relative investment prices) increases the share of firms experiencing spikes by about 2 percentage points and raises TFP by approximately 1% one year later. These magnitudes can calibrate or validate models of embodied technological change.

Interestingly, the TFP response exceeds what standard neoclassical growth theory predicts. The mechanical effect of investment-specific shocks on measured TFP should be approximately equal to the size of the shock times the share of capital in production. The amplified response suggests that embodied technology adoption triggers additional productivity-

enhancing adjustments. These findings highlight that capital-embodied technological change operates through richer mechanisms than simple factor accumulation, with important implications for understanding aggregate productivity dynamics.

The estimated TFP response merits further discussion, as it substantially exceeds benchmark predictions. In a standard neoclassical growth model, the response of measured TFP to investment-specific productivity shocks equals capital's share in production times the shock size. With an estimated capital share of 0.18 (reflecting tangible capital in our Italian manufacturing sample), the predicted TFP effect is approximately 0.13% ($0.18 \times 0.7\%$, the standard deviation of the shock)—nearly eight times smaller than our estimated 1.0% response. This substantial amplification suggests that embodied technology adoption triggers productivity-enhancing adjustments beyond simple factor accumulation.

This result motivates our theoretical framework. If investment-specific shocks generate such substantial effects at the firm and sectoral levels, how do these micro-level responses aggregate to shape macroeconomic dynamics? Does firm heterogeneity in technology adoption amplify or dampen aggregate volatility? To address these questions, we develop a general equilibrium model in Section 4 that explicitly incorporates the microeconomic features documented in our empirical analysis: lumpy capital adjustment, heterogeneous investment timing across firms, and the negative relationship between investment age and productivity. The model allows us to trace how individual firms' technology adoption decisions aggregate through capital accumulation to generate macroeconomic fluctuations, and to assess whether the amplification mechanisms evident in our sectoral estimates materially affect aggregate outcomes.

4 Model: Firm Heterogeneity and Technology Adoption

We formulate a general equilibrium framework featuring firm heterogeneity and investment-specific technological progress to study the aggregate implications of investment age heterogeneity.

Our model builds on the lumpy investment framework of Khan and Thomas (2003, 2008), incorporating three key features. First, following Solow (1960) and Greenwood, Hercowitz and Krusell (1997), technological progress is embodied in investment goods—newer vin-

tages are more efficient at converting physical capital into productive services. Second, departing from Khan–Thomas and standard vintage models, firms can choose to invest in capital goods embodying different levels of efficiency: the latest vintage (most efficient) or their current vintage (less efficient). Third, adopting the latest vintage requires paying a non-convex switching cost, creating an (S, s) decision rule: firms adopt when benefits sufficiently exceed costs, otherwise postponing and continuing with their current vintage.

These features generate a non-degenerate equilibrium distribution of technologies across firms. Unlike neoclassical growth models where all firms instantaneously adopt the technological frontier, our framework features heterogeneity in technology efficiency: some firms operate at the frontier while others operate older vintages. Aggregate and idiosyncratic shocks alter optimal adoption timing, shifting the cross-sectional distribution of technologies and generating endogenous productivity movements beyond the direct shock impact that lead to amplified responses in output.

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

4.1 Production

The economy consists of a continuum of firms normalized to one. As common in the literature on firm heterogeneity and investment-specific technological progress, each firm has access to an increasing and concave production function (F) that combines a predetermined stock $\tilde{k}$ of capital services rather than physical units and labor hired on a spot market to produce output y. We report model variables around the balanced growth path, deflated by their respective deterministic trends.[8] For better readability, we drop individual firms subscripts; the production function is:

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

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

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

4.2 Firm's Technology Adoption and Investment Decision

The main innovation of the framework consists of modelling a vintage technology structure for capital-embodied technological change. The stock of capital *services*, $\tilde{k}$, evolves according to the following

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

where $\delta \in (0, 1)$ is the rate of physical capital depreciation and γ_k is the deterministic gross growth rate of capital services around the balanced growth path. As in Solow (1960) and Greenwood, Hercowitz and Krusell (1997, 2000), we formalize the notion of investment-specific technological change by assuming that the inflow of new capital services, $\widetilde{i}_{j,t}$, is the product of the physical quantity of investment good, $i_{j,t}$, and its associated vintage productivity, denoted with a slight abuse of notation by $q_{j,t}$.[9] Alternatively, $1/q_{j,t}$ can be interpreted as representing the cost of producing one unit of capital (in terms of final output) of a given vintage j. The index j takes nonnegative values and identifies the current vintage of technology available to each firm and is optimally chosen by the firm, with $j = 0$ denoting the latest vintage and higher values of j indicating older vintages.

Note that $\tilde{k}$ denotes the stock of *capital services* obtained with investment goods and therefore does not have a subscript j. As in Khan and Thomas (2003) and in the spirit of Solow (1960) and the subsequent literature, expressing the stock of capital services in units of efficiency rather than physical units allows for a convenient aggregation of physical units of investment embodying different levels of efficiencies, accumulated by the firm over time.

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

The investment-specific technological frontier grows deterministically at the gross rate of $\gamma_q > 1$. Around the balanced growth path, the efficiency of the latest vintage $q_{0,t}$ fluctuates around its deterministic trend and follows a Markov process in logs. Similarly to ε, $q_{0,t}$ is discretized so that in every period $q_0 \in \{q_{0,1}, q_{0,2}, ..., q_{0,N_q}\}$, where $Pr(q_0 = q_{0,z} | q_0 = q_{0,n}) \equiv \pi_{nz}^{q_0} \geq 0$, and $\sum_{z=1}^{N_{q_0}} \pi_{nz}^{q_0} = 1$ for each $n = 1, ..., N_{q_0}$.

The timing of events is as follows. Each firm starts the period with a predetermined stock of capital services, $\tilde{k}_t$, and a given vintage, $q_{j,t}$. The idiosyncratic productivity ε_t and the efficiency of the capital-embodied technology $q_{0,t}$ are realized, and firms draw the identically and independently distributed cost of adopting the latest vintage ξ. Then, firms decide whether to pay the cost to switch technology by acquiring the latest investment good and choosing the capital stock for the next period. Each unit of capital acquired comes at a cost of $1/q_{0,t}$. A firm that pays the adjustment cost adopts the latest vintage in the current period so that $\tilde{k}_{t+1}$ depends on the efficiency $q_{0,t}$ and the quantity of investment $i_{0,t}$ undertaken. A firm that postpones adopting the latest vintage chooses $\tilde{k}_{t+1}$. In this case, the firm keeps acquiring investment goods embodying a technology that becomes less efficient relative to the deterministic trend at a rate γ_q ($q_{j,t}/\gamma_q$). At the beginning of the next period, the efficiency of the adopters of the latest vintage is $e^{\rho_q \times log(q_{0,t})}$ as $q_{0,t}$ reverts to its mean at a rate of ρ_q. Similarly, firms that have postponed adjusting start the next period with a vintage that has efficiency equal to $q_{j,t}/\gamma_q$. To be clear, firms not adopting the latest vintage keep acquiring investment goods embodying a constant *level* of efficiency, but, around the balanced growth path and relative to the frontier, become less efficient given the availability of newer investment goods embodying a more efficient technology. In the next section, we describe the tradeoffs associated with the choice of technology adoption.

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

In each period, a firm is defined by its vintage productivity q_j, its idiosyncratic productivity level $\varepsilon \in \mathcal{E} \equiv \{\varepsilon_1, \varepsilon_2, ..., \varepsilon_{N_\varepsilon}\}$, its predetermined stock of capital $\tilde{k}_t \in \mathbf{R}_+$, and its cost associated with vintage adoption $\xi \in [\underline{\xi}_l, \overline{\xi}_h]$, which is denominated in units of labor and drawn from the time-invariant distribution G common to all production units and assumed to be uniform. As the firm's current adjustment cost has no implications for its future adjustment, the distribution of firms is summarized by $(\varepsilon_t, q_{j,t}, \tilde{k}_t)$. To characterize the distribution of firms over $(\varepsilon_t, q_{j,t}, \tilde{k}_t)$, we use the probability measure μ defined on the Borel algebra S for the product space $S = \mathcal{E} \times \mathbf{R}_+ \times \mathbf{R}_+$. The aggregate state of the economy is described by $(q_{0,t}, \mu_t)$, the efficiency of the technology frontier and the distribution of firms that evolves according to a mapping $\Gamma : \mu_{t+1} = \Gamma(q_{0,t}, \mu_t)$.

Three features distinguish our framework from existing models. First, unlike standard vintage capital models (Solow, 1960; Greenwood, Hercowitz and Krusell, 1997) where all firms invest in the latest vintage, we allow firms to postpone adoption and continue investing in older vintages. This generates endogenous heterogeneity in technology *efficiency*, not just *quantity*. Second, unlike Khan and Thomas (2003, 2008) where all firms access identical investment goods, investment goods in our model embody different efficiencies based on the firm's adoption history. Third, unlike models with exogenous TFP heterogeneity (Hopenhayn, 1992), productivity differences in our model partly reflect endogenous technology choices. When the adoption cost $\bar{\xi}_{l,h} = 0$, our model collapses to a standard neoclassical growth model with investment-specific technical progress where all firms operate at the frontier.

4.3 Firm's Dynamic Programming Problem

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

Let $V^1(\varepsilon_t, q_{j,t}, \tilde{k}_t, \xi_t; q_{0,t}, \mu_t)$ denote the expected discounted value of a firm entering the

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

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

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

$$\begin{aligned} V^1(\varepsilon_t, q_{j,t}, \tilde{k}_t, \xi_t; q_{0,t}, \mu_t) = \max_{n_t, \tilde{k}^A_{t+1}, \tilde{k}^{NA}_{t+1}} & \underbrace{\left[\varepsilon_t F(\tilde{k}_t, n_t) - w(q_{0,t}, \mu_t) n_t\right] p(q_{0,t}, \mu_t)}_{\text{Current period profits}} \\ & + \max \left\{ \begin{array}{l} \text{Value of Adoption:} \\ \dfrac{(1-\delta)\tilde{k}_t}{q_{0,t}} p(q_{0,t}, \mu_t) - \xi_t w(q_{0,t}, \mu_t) + R(\varepsilon_t, q_{0,t}, \tilde{k}^A_{t+1}; q_{0,t}, \mu_t), \\ \text{Value of No adoption:} \\ \dfrac{(1-\delta)\tilde{k}_t}{q_{j,t}/(\gamma_q q_{0,t})} p(q_{0,t}, \mu_t) + R(\varepsilon_t, q_{j,t}, \tilde{k}^{NA}_{t+1}; q_{0,t}, \mu_t) \end{array} \right\} \end{aligned} \tag{5}$$

$$\text{subject to: } n_t, \tilde{k}^A_{t+1}, \tilde{k}^{NA}_{t+1} \in \mathbb{R}_+$$

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

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

where firms buy capital goods that embody the latest vintage $q_{0,t}$ at time t and have a continuation value at the beginning of the period evaluated over the possible realizations of the

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

4.4 Households

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

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

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

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

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

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

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

and $\tilde{k}_{t+1}$ capital services.

4.5 Recursive Equilibrium

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

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

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

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

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

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

(vi) Model-consistent dynamics: The evolution of the cross-sectional distribution that characterizes the economy, $\mu_{t+1} = \Gamma(q_{0,t}, \mu_t)$, is induced by the adjustment decision and the exogenous processes for ε. Conditions (i), (ii), (iii), (iv), and (v) define an equilibrium given Γ, while condition (vi) determines the equilibrium condition for Γ. We confine to Appendix E 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 Capital-Embodied Technology

In this section, we discuss the role of investment-specific productivity heterogeneity for aggregate dynamics.

At the firm level, non-convex adoption costs imply that technology adoption and investment decisions follow an (S, s) rule: some firms invest in capital goods embodying the

latest vintage technology, while others postpone adoption and continue investing in goods embodying their current vintage.[10] Because the latest vintage's efficiency grows over time at a positive rate, the efficiency gap between the frontier and other vintages widens. For each firm, this gap typically exceeds the size of any individual shock, as it reflects both accumulated deterministic drift from the frontier and current innovations. The efficiency gap, combined with realizations of idiosyncratic and aggregate shocks, determines the timing of technology adoption at the firm level. Consequently, firms invest in vintages with different efficiency levels. Shocks to frontier technology $q_{0,t}$ alter the timing of adoption: when firms upgrade from outdated vintages, they generate persistent efficiency gains that exceed the direct shock impact. Once a firm moves to the frontier, even accounting for the drift in the evolution of the frontier (γ_q), its efficiency $q_{j,t}$ will, on average, be higher than its steady-state value.

At the aggregate level, changes in $q_{0,t}$ translate into aggregate fluctuations through changes in the cross-sectional distribution of firm adoption choices. These individual decisions affect aggregate capital services and output through both a *quantity* and an *efficiency* channel, jointly determining total physical investment ($i_{j,t}$) and the *average* efficiency embodied in those investment goods ($q_{j,t}$). For example, when the latest vintage's efficiency exceeds expectations, more firms pay the fixed cost to adopt it. This increases aggregate capital services through higher total investment (quantity channel) and higher average efficiency of that investment (efficiency channel), since the optimal capital stock increases with the vintage quality operated by the firm. As adopting firms continue operating technologies near-frontier vintages in periods after the shock, the average efficiency of investment remains elevated above its steady-state level for several periods, creating persistence in average efficiency even after $q_{0,t}$ reverts toward its mean. This persistent compositional effect propagates the shock's impact through time. In sum, the interaction of the quantity and the efficiency channels contributes to the amplification of aggregate dynamics relative to the neoclassical growth model with zero nonconvex costs, as we discuss in Section 7.

This compositional mechanism contrasts sharply with Khan and Thomas (2003), where firm heterogeneity proves irrelevant for aggregate dynamics. In their framework, all firms access identical investment goods, so heterogeneity reflects capital *quantity* (firms have dif-

[10] See Appendix E for additional details.

ferent capital stocks) but not *efficiency* (all capital has the same productivity). General equilibrium forces—particularly real interest rate adjustments—neutralize the impact of adjustment cost heterogeneity on aggregate investment demand.

Before describing the calibration strategy of the model, it is worth examining the role of three key parameters. The support of the adjustment cost distribution $(\overline{\xi}_l, \overline{\xi}_h)$ determines the magnitude of the cost of adjusting technology. The higher value of the adjustment cost, the higher the potential cost of adopting the latest vintage. Increasing this cost leads to a higher average investment age increasing the mass of firms experiencing spikes. The idiosyncratic process's persistence and standard deviation interact with the vintage effect in shaping the economic incentives that make the firm adopt the latest vintage and choose the appropriate level of capital. We now turn to Section 5 where we calibrate these parameters to match microeconomic data on investment behavior and TFP.

5 Taking the Model to the Data

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

5.1 Parameterization

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

Parameter values based on a priori information. One period in the model represents

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

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

Table 3: Benchmark Calibration

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

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 instruc-

tive to describe the calibration as a few distinct steps. The lower and the upper support of the technology adjustment cost function ($\overline{\xi}_l$,$\overline{\xi}_h$) are set to reproduce the share of positive investment accounted for by firms experiencing spikes and the average investment age of the firms' cross-sectional distribution, respectively. The estimated $\overline{\xi}_l$ and $\overline{\xi}_h$ imply that adjustment costs, conditional on adjustment, are equal to about 3% of annual output, which is on the lower end of the previous estimates in the existing literature; see, for instance, Bachmann, Caballero and Engel (2013).[11] The persistence of the exogenous idiosyncratic productivity process (ρ_ε) is set so that the model firms' TFP displays the same persistence as the one estimated in the data (0.80). This coefficient obtains fitting an autoregressive process of order one to the logarithm of firm-level TFP.[12] The standard deviation of the idiosyncratic productivity process (σ_ε) is selected to reproduce the fraction of firms experiencing an investment spike in the data. In Section 7, we compare the business cycle performance of the vintage model with a standard neoclassical growth benchmark with firm heterogeneity. Our vintage model nests the neoclassical benchmark as a special case that is obtained when the technology adjustment cost is equal to zero and there is no technological regress.

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.[13] As a result, we proxy the distribution with the mean capital stock, $\overline{K}_{t+1}$, and the mean of the investment-specific productivity distribution, $\overline{Q}_{t+1}$, and their interaction. We use the same set of regressors to forecast the marginal utility of consumption, p_t. Moreover, we estimate the rule conditional on each realization of the aggregate

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

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

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

process, $q_{0,t}$. We report additional details regarding the solution method in Appendix F.1.

6 Microeconomic Heterogeneity in the Model and the Data

We now discuss the ability of the model to fit multiple dimensions of the microeconomic heterogeneity in the data related to firms' investment behavior and TFP. To ensure a consistent comparison of the model-generated statistics with the data, we follow the same measurement of capital and TFP described in Section 2.2, using the perpetual inventory method and physical depreciation to construct the physical capital stock. In Sections 6.1 and 6.2, we show that the model fits well the cross-sectional distribution of investment rates and the timing of investment spikes across firms—i.e., the empirical proxy for vintage technology. In Section 6.3, we discuss the relationship between dispersion in idiosyncratic investment-specific productivity in the data and in the model. In Section 6.4, we interpret the empirical relationship between investment age and TFP through the lens of the model-based simulated series.

6.1 Cross-Sectional Distribution of Investment Rates

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

The average share of firms experiencing a positive spike is matched by virtue of the calibration strategy discussed in Section 5.1. The model closely matches the fraction of firms experiencing negative spikes and positive investment. However, the model overstates the share of firms downsizing capital and understates the share of inactive firms. A comparison with a neoclassical growth benchmark sheds light on the role of vintage technology

Table 4: Distribution of Firm Investment Rates

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

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

in driving aggregate dynamics. Specifically, we consider the model "Neoclassical $\overline{\xi}_{l,h} = 0$" that retains the same set of parameters as the vintage model and sets the adjustment cost to technology, $\overline{\xi}_l$ and $\overline{\xi}_h$ equal to zero.

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

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

6.2 Cross-Sectional Distribution of Investment Age

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

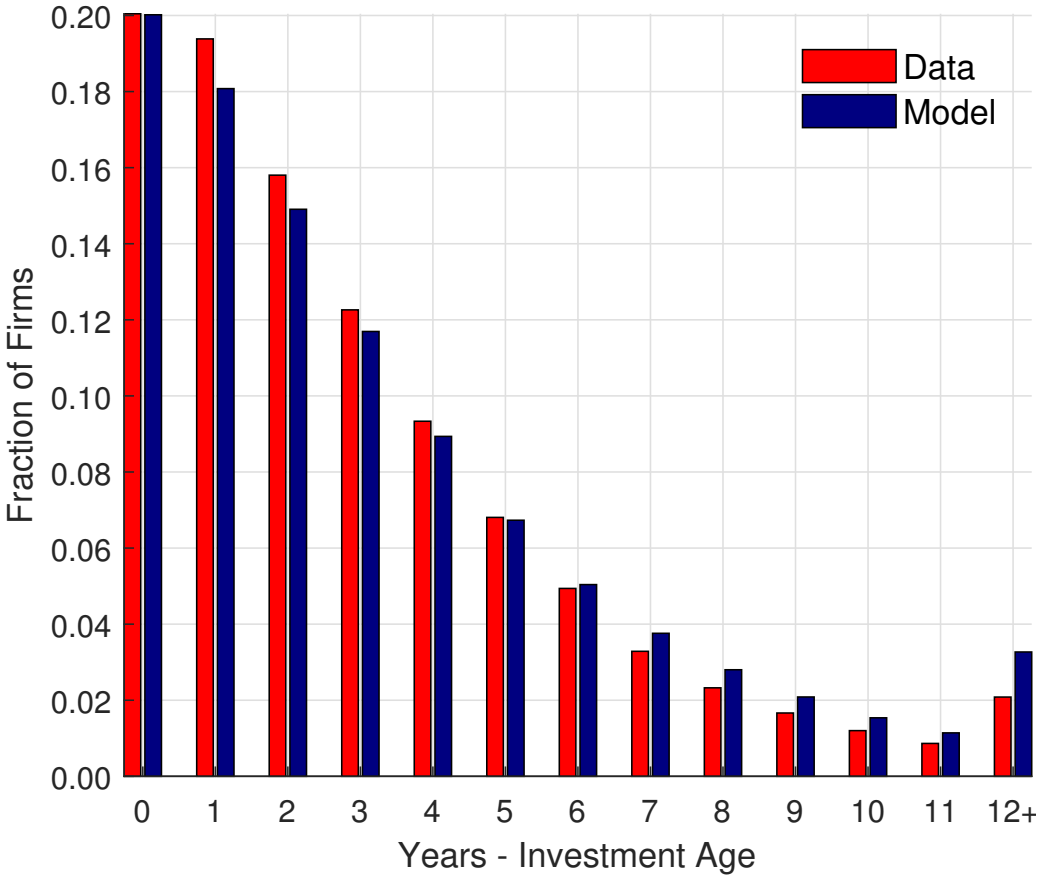


Figure 2: Comparison between the Empirical and the Model-Based Investment Age Distribution

We remind the reader that our calibration targets the average investment age of about 3 and a half years and, as discussed in the previous subsection, the fraction of firms experiencing an investment spike—i.e., firms exhibiting an investment age equal to zero. As shown in Figure 2 the model distribution closely matches the data, with a similar decay in investment age. Overall, the performance of the model is satisfactory. We notice that the ability of the framework to reproduce the timing of investment spikes across firms is distinct from

the model's success in accounting for the cross-section of investment rates. While the fraction of firms with an investment age of zero coincides by construction with the fraction of firms exhibiting spikes, the fraction of firms with an age of one and above depends on the investment behavior of firms—determined by the realization of individual states.

For comparison, setting nonconvex adjustment costs of technology adoption equal to zero, reduces the average investment age of the model to about 3 years.

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

We now examine how heterogeneity in investment-specific productivity contributes to dispersion in *TFP* across firms. As discussed in Greenwood, Hercowitz and Krusell (1997) and in Section 2.2, we measure the capital stock in terms of consumption (rather than units of efficiency), implying that productivity gains embodied in newer, more efficient capital are thus absorbed into measured TFP.[14] Using the Solow residual, we compute TFP for each firm and the resulting cross-sectional dispersion in the vintage model and in the neoclassical growth benchmark with adoption cost $\overline{\xi}$ equal to zero. To assess the contribution of vintage technology to the cross-sectional dispersion, we compute the log of the 90/10 range of firm-level TFP in the model and compare it with Cerved data. The latter obtains as dispersion of residuals from regressing the log of TFP on sector, time, and sector-time effects over the 1998-2016 sample, the same specification used in the empirical analysis in Section 3.

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

[14] Greenwood, Hercowitz and Krusell (1997) show that the neoclassical growth model with investment-specific technical progress can be represented by a model with neutral technological change. In this case, the capital stock is measured in terms of consumption (rather than units of efficiency), and the relative price is constant and equal to one. Under this transformation, the measured Solow residual combines neutral and investment-specific productivity; see, for a discussion, Appendix B on page 360 in Greenwood, Hercowitz and Krusell (1997).

Table 5: Data and Model: TFP Dispersion

	Data	Baseline Vintage	Neoclassical $\overline{\xi}_{l,h} = 0$
	(A)	(B)	(C)
$\log(TFP_{90}/TFP_{10})$	1.00	0.08	0.06
Decomposition			
$\log(q_{j,90}/q_{j,10})$	n.a.	0.09	0
$\log(\varepsilon_{90}/\varepsilon_{10})$	n.a.	0.06	0.06

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

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

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

6.4 Model and Data: TFP and Investment Age

Having established that our framework accurately matches the cross-sectional distribution of investment rates and investment age, we now use model-generated data to support the

interpretation that the negative relationship between TFP and lagged investment age documented in Section 3 reflects investment-specific technical change.

We first establish that the negative relationship between TFP and investment age cannot be replicated by a standard lumpy investment model such as Khan and Thomas (2008) or Bachmann, Caballero and Engel (2013). Using the Bachmann, Caballero and Engel (2013) model—which features idiosyncratic and aggregate neutral TFP shocks and nonconvex adjustment costs—we simulate a panel of firms and estimate equation 1 on the artificially generated data.[15] This specification projects current TFP on the first lag of investment age while controlling for lagged TFP. Controlling for lagged TFP ensures that current TFP is orthogonal to any variable dated t-1. Consequently, the estimated coefficient on lagged investment age is zero in this model. The coefficient is zero because in models with only neutral TFP shocks, lagged investment age contains no information about embodied technology and thus has no predictive power for current TFP once lagged TFP is controlled for.

By contrast, in models with investment-specific technical change—our vintage capital model and its neoclassical growth counterpart with zero adoption costs—the same exercise yields a negative and statistically significant coefficient on lagged investment age. Table 6 shows that both models with vintage technology display a negative relationship between TFP and the lag of investment age. Moreover, the estimated coefficient in the vintage model with positive adoption costs is about five times larger in absolute value than the coefficient in the neoclassical growth model with zero adoption costs ($\xi_{l,h} = 0$), bringing it closer to the empirical estimate. This difference reflects the larger productivity gains that arise in the vintage model when firms upgrade from obsolete technology to the frontier, contrasting sharply with the neoclassical growth model where firms continuously operate at the technological frontier. Together with the structural interpretation in Section 3.2 using investment-specific shocks identified in the SVAR, the evidence in Table 6 supports investment-specific technical progress as the force behind the negative empirical relationship documented in Section 3.

[15] We simulate data for a sample of 3000 firms, enough points to replicate the cross-sectional distribution of investment rates, for 5000 periods. We take the latest 70 observations to estimate equation 1.

Table 6: Data and Model: TFP and Investment Age

	$TFP_{f,t}$ on $\frac{y_{j,t}}{k^{\theta}_{j,t-1} n^{v}_{j,t}}$	$TFP_{f,t}$ on $q^{\theta}_{j,t-1}$	$TFP_{f,t}$ on ε_t
	(A)	(B)	(C)
Data			
$Inv.Age_{f,t-1}$	-0.335	n.a.	n.a.
Investment-Specific Technological Progress			
Baseline Vintage			
$Inv.Age_{f,t-1}$	-0.062	-0.121	-0.011
Neoclassical Growth $\overline{\xi}_{l,h} = 0$			
$Inv.Age_{f,t-1}$	-0.012	n.a.	-0.012
Neutral Technological Progress			
Lumpy Investment $\overline{\xi}_{l,h} > 0$			
$Inv.Age_{f,t-1}$	-0.000	n.a.	-0.000
Neoclassical Growth $\overline{\xi}_{l,h} = 0$			
$Inv.Age_{f,t-1}$	-0.000	n.a.	-0.000

Note: 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% at time t-1.

7 Vintage Technology Amplifies Aggregate Fluctuations

In this section, we demonstrate that microeconomic heterogeneity in investment-specific productivity amplifies business cycle dynamics in response to investment-specific technology shocks, relative to a neoclassical growth benchmark. To understand the propagation mechanism, in Section 7.1, we report impulse response functions for both aggregate variables and cross-sectional distributions. Section 7.2 examines the business cycle properties of aggregate series generated from the model. Finally, Section 7.3 positions our contribution within the existing literature on lumpy investment.

7.1 Inspecting the Mechanism: Impulse Response Functions

In this section, we illustrate the mechanism through which investment-specific technology shocks propagate in the vintage model, generating amplified aggregate dynamics relative to the neoclassical growth model. The vintage model substantially amplifies the propagation of investment-specific technology shocks. Following identical shocks, the GDP response peaks at 1% with greater persistence, compared to approximately 0.3% in the neoclassical benchmark. To examine these quantitative differences in propagation mechanisms, we compute impulse response functions by tracing the dynamics following a one-standard-deviation shock to the efficiency of the latest vintage ($q_{0,t}$), with both economies initially at their respective stochastic steady states. We obtain these steady states by simulating the models forward under the assumption that $q_{0,t}$ remains constant at its mean value.

Neoclassical Growth Benchmark Dynamics. Figure 3 contrasts the dynamics of the neoclassical benchmark (dashed red lines) with those of the vintage model assuming that the discretized exogenous productivity process for the efficiency of the latest vintage follows an annual persistence of 0.65. As is standard in the neoclassical growth model, firms respond on impact by increasing both investment and employment, since the latest vintage of investment goods embodies more efficient technology; see, for instance, Greenwood, Hercowitz and Krusell (2000). This combination of higher investment and improved technology raises aggregate capital services one period after the shock. Crucially, zero adoption costs in the neoclassical growth model ensure that all firms immediately invest in goods embodying the latest vintage. Consequently, the average efficiency of investment goods instantaneously equals the efficiency of the latest vintage. Because of the time-to-build assumption in capital accumulation, output is initially unaffected by the improved efficiency of investment goods. Households therefore finance higher investment by reducing consumption on impact. Over time, as more efficient capital is installed in production, higher capital services increase output, allowing consumption to recover. As the shock dissipates, the economy gradually returns to its stochastic steady state.

Vintage Model: Amplification and Vintage Heterogeneity. We now examine how the vintage capital model generates amplified aggregate dynamics compared to the standard neoclassical growth model. The mechanism underlying this amplification centers on the

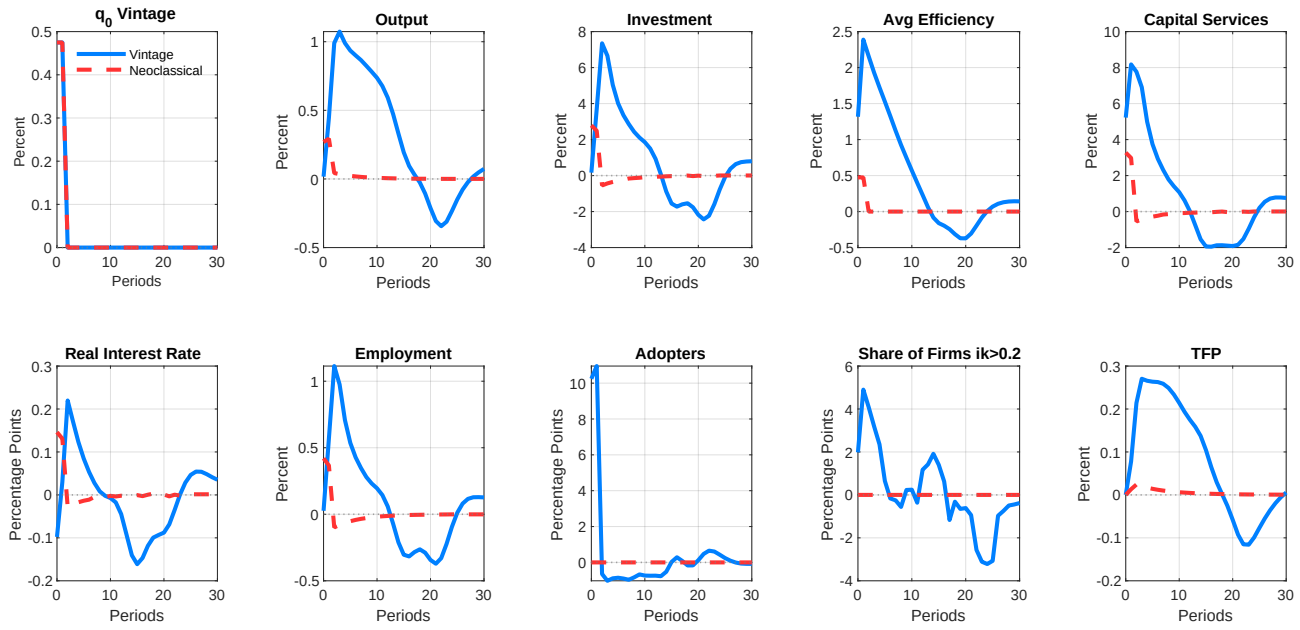


Figure 3: Impulse response functions following an increase in the efficiency of the latest vintage q_0 in the vintage model (solid line) and in the neoclassical growth benchmark with $\overline{\xi}_{l,h} = 0$ (dashed line).

response of capital services, a composite measure that incorporates both the quantity of physical capital and its average efficiency. In the vintage framework, capital services expand through two channels: physical units of investment and efficiency gains as firms transition from older, less productive technologies to newer, more efficient ones.

The average efficiency margin operates differently across the two models. In the neoclassical growth framework, as all firms invest in the latest technology, changes in average investment efficiency reflect only the direct impact of the productivity shock itself. In contrast, the vintage model features an additional, endogenous amplification mechanism: as more firms adopt newer technologies following a positive shock, the composition of investment shifts toward more efficient vintages, raising average efficiency beyond what the shock alone would generate. This compositional effect arises because the share of firms operating newer technologies (adopters) rises above its steady-state level, as firms switch from their previous technology $q_{j,t}$—which has fallen behind the technology frontier at rate γ_q for each period the firm has not invested—to the frontier technology $q_{0,t}$. The physical quantity of investment also increases, as the higher efficiency of investment goods stimulates capital formation. This same mechanism leads to a persistent response of productivity, investment, and output beyond the one implied by the initial shock. Because the share of

high-productivity firms remains elevated above its steady-state level, output exhibits a persistent increase that exceeds what the initial shock alone would produce. The higher average efficiency of investment results in higher TFP relative to the neoclassical benchmark.

Having outlined the core mechanism, we now analyze the impulse responses shown in Figure 3 to examine the timing of effects. On impact, the shock directly raises the efficiency of new investment goods. However, due to the standard time-to-build assumption in capital accumulation, these investment goods do not immediately enter production. Consequently, the initial output response is driven primarily by employment adjustments. In the vintage model, the amplified efficiency gains permit households to achieve their desired capital accumulation plan with less physical investment, resulting in a smaller decline in consumption and a correspondingly smaller increase in employment (and thus output) on impact.

In subsequent periods, newly installed capital enters production, generating a substantial boost in output. Investment continues to rise as two reinforcing factors sustain the boom: first, average efficiency remains above its steady-state level; second, the share of high-productivity firms remains elevated above its steady-state level stimulating investment. The compositional effect thus generates a more persistent response of TFP and aggregate variables than in the neoclassical growth model, where the shock's impact dissipates more quickly.

Because of this amplified response of average efficiency, capital services increase more substantially in the vintage model than in its neoclassical counterpart, where only the quantity margin operates. As shown in the figure, capital services increase by about 8% in the vintage model, approximately twice as much as in the neoclassical growth model. This stronger capital services response is achieved despite a smaller initial increase in physical investment and a correspondingly smaller drop in consumption relative to the neoclassical growth model. As capital services become productive under the time-to-build constraint, output continues to expand in subsequent periods, with the vintage model exhibiting both higher peak responses and greater persistence. We also note that the reversion to the steady state is not monotonic as in the neoclassical model. Instead, the surge in investment following the shocks displays echo effects over time, similar to those documented in machine replacement models such as Cooper and Haltiwanger (1993). These cyclical patterns arise because the initial wave of investment creates synchronized replacement needs when the

capital stock depreciates, generating peaks in investment demand that gradually dampen as the economy returns to steady state.

We now turn to the role of general equilibrium effects, specifically the responses of wages and the real interest rate. As in the classic Khan-Thomas framework, general equilibrium forces work against our amplification mechanism: market-clearing prices act as a stabilizing force, dampening aggregate fluctuations. Appendix G quantifies this tension, showing that endogenous price adjustments reduce the volatility of output and investment by approximately 80% and 90%, respectively. Yet despite these offsetting forces, our vintage technology channel generates substantial amplification. Notably, the model's predicted real interest rate dynamics—both in sign and magnitude—align well with the empirical patterns we identify in the SVAR analysis (Appendix D), providing independent validation of the model's general equilibrium properties in response to investment-specific shocks.

Quantitative Responses in the Data and in the Models. We now compare the quantitative model's predictions with the estimated effects from Section 3.2 regarding changes in investment age, the share of firms experiencing investment spikes, and TFP response. The share of firms experiencing investment spikes increases by about 2 percentage points in the model, closely matching the increased observed in the data. The model further reveals that this increase is driven by cyclical variation in technology adoption rather than idiosyncratic factors. Specifically, the share of all firms experiencing investment spikes increases by approximately 2 percentage points, driven entirely by firms adopting the latest vintage—embodying frontier technology. In contrast, the share of firms experiencing spikes without adopting the latest technology remains relatively constant over the business cycle.

The average efficiency of investment goods manifests as total factor productivity when physical capital stock, rather than capital services, serves as the production input.[16] The vintage model generates a TFP increase of about 0.3% at the peak, compared to about 0.8% in the data—substantially closer than the neoclassical model's 0.02% response. Overall, the vintage model successfully replicates the three key empirical patterns, providing quantitative support for the compositional channel identified in Section 3.2.

[16]We define aggregate TFP using a standard aggregate production function. Because returns to scale are close to unity, computing TFP at the firm level and then aggregating yields nearly identical results to using an aggregate production function directly.

7.2 Business Cycle Properties

We now compare the business cycle properties of HP-filtered aggregate series obtained by simulating the baseline vintage model with a neoclassical growth benchmark—that is, the vintage model with zero adoption costs, $\overline{\xi_l}, \overline{\xi_h} = 0$.

We have already discussed the propagation mechanism in the previous section; here we highlight two key findings. First, even when filtering the data to examine behavior at business cycle frequencies, the vintage model displays higher volatility of aggregate series relative to the neoclassical benchmark when both economies are driven by identical shocks. As discussed in the impulse response analysis, this amplification depends on the joint, reinforcing effect of investment quantity and the associated improvement in average efficiency.

Specifically, aggregate shocks alter firms' incentives to adopt new technologies, shifting the cross-sectional distribution of investment-specific productivity and generating aggregate fluctuations in the efficiency of investment and TFP beyond the initial impulse from the stochastic technological frontier. The larger responses of more efficient investment and employment, in turn, translate into stronger output dynamics.

Second, the vintage model delivers greater persistence in cyclical series than the standard neoclassical growth model due to heterogeneity in vintage technology, improves the endogenous propagation of technology shocks. Unlike the neoclassical model where shocks dissipate quickly, the vintage framework generates persistent effects through the share of high-productivity firms that remains elevated above its steady-state level.

It is now well understood in the literature that a single shock, even when plausibly calibrated, cannot account for entire business cycle dynamics. In the vintage model, calibrated investment-specific shocks account for approximately one-third of GDP volatility in the data and about 60% of cyclical investment fluctuations, results consistent with Smets and Wouters (2007) and Justiniano, Primiceri and Tambalotti (2010).[17]

In Appendix H, we demonstrate that even when examining simulated raw series, the vintage model displays higher volatility and persistence relative to the standard neoclassical growth model. Specifically, Table A.8 reports business cycle statistics computed on the stationary variables of the model, as both models share the same deterministic trend.

[17] The remaining GDP volatility is accounted for by other shocks, such as those to monetary or fiscal policy, uncertainty, or competition, among others.

Table 7: Business Cycle Statistics: Technology Shock

	GDP	Consumption	Investment	Employment
	(A)	(B)	(C)	(D)
Data				
σ_X	2.10	1.25	6.65	1.08
$Corr(X_t, X_{t-1})$	0.68	0.56	0.54	0.71
$Corr(X_t, GDP_t)$	1	0.88	0.91	0.81
Baseline Vintage				
σ_X	0.56	0.32	3.60	0.57
$Corr(X_t, X_{t-1})$	0.89	0.88	0.81	0.80
$Corr(X_t, GDP_t)$	1	0.24	0.88	0.84
Neoclassical model $\overline{\xi_l}, \overline{\xi_h} = 0$				
σ_X	0.34	0.24	3.24	0.55
$Corr(X_t, X_{t-1})$	0.57	0.53	0.41	0.40
$Corr(X_t, GDP_t)$	1	-0.39	0.90	0.88

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

7.3 Relationship with the Lumpy Investment Literature

Our contribution to the existing literature is to provide a theoretical mechanism highlighting how firm heterogeneity shapes aggregate dynamics through investment-specific technological change, empirically grounded by the evidence presented in Section 3. The existing literature on lumpy investment has extensively debated the role of firm heterogeneity in models with exogenous TFP, particularly regarding the relative importance of partial versus general equilibrium forces. This debate is exemplified by a series of influential papers by Caballero and Engel, later joined by Bachmann, and by Khan and Thomas. Caballero and Engel (1999) and Bachmann, Caballero and Engel (2013) argued in favor of models with lumpy investment to explain the procyclical sensitivity of US aggregate investment to shocks. However, Thomas (2002) and Khan and Thomas (2003, 2008) emphasized that general equilibrium effects can significantly dampen—and nearly erase—the partial equilibrium forces arising from microeconomic heterogeneity, yielding the so-called irrelevance result. Building on this literature, Winberry (2021) and Koby and Wolf (2019) provided important clarity by linking these differences to the partial equilibrium elasticity of investment demand with

respect to the real interest rate. Winberry showed that the Khan and Thomas (2008) framework implies an elasticity exceeding 1,000—substantially higher than estimates from fiscal policy studies—and demonstrated that an alternative calibration with a more plausible elasticity delivers approximately 10 and 20% amplification in aggregate output and investment dynamics, respectively, while preserving general equilibrium effects.

Our framework complements and extends this literature by introducing a distinct amplification mechanism based on heterogeneity in investment goods efficiency. Our model with vintage capital amplifies aggregate investment-specific shocks while featuring an elasticity of investment with respect to the real interest rate of -305, orders of magnitude larger than Winberry (2021) and Koby and Wolf (2019). Importantly, we emphasize an additional channel operating through investment-specific technological change and endogenous technology adoption via capital accumulation. The key insight is that heterogeneity in investment-specific efficiency generates time-varying changes in the average efficiency of investment that exceed the gains directly implied by the productivity shock itself. As discussed in Section 7.1, compositional changes in the cross-sectional distribution of embodied technology—arising from procyclical shifts in adoption patterns across firms—amplify aggregate output dynamics by approximately 50% relative to the standard neoclassical growth model.

Unlike the literature cited above, the source of this amplification lies in the endogenous evolution of the technology distribution. As detailed in Section 7.1, widespread adoption by laggard firms creates procyclical shifts in the technology distribution that amplify movements in the aggregate production possibility frontier beyond what occurs in the neoclassical growth model. The amplification arises because adoption is procyclical—more firms upgrade during booms—creating an endogenous boost to the average efficiency of investment beyond that implied by the investment-specific shock alone, an effect that is absent when all firms access the frontier. This contrasts with frameworks emphasizing investment-specific technical change, such as Solow (1960), Greenwood, Hercowitz and Krusell (1997), and Khan and Thomas (2003), where the absence of heterogeneity in investment efficiency implies that all firms instantaneously adopt frontier technology, eliminating this compositional channel. Consequently, our mechanism delivers amplification through a fundamentally different channel—the endogenous evolution of the technology distribution—within a microeconomically founded general equilibrium framework.

Our model features technology adoption through investment and embodied technology, complementing models that focus on the innovation side. Recently, Wang and Zhang (2025) propose a model of lumpy R&D building on Comin and Gertler (2006). As is common in endogenous growth models such as Romer (1990), in Wang and Zhang the quality and quantity of ideas are isomorphic: raw ideas and R&D ideas differ only in whether they are actively pursued, and all ideas are treated as homogeneous once processed through R&D projects. This isomorphism eliminates any meaningful distinction between more versus better technology. Our model fundamentally departs from this approach by treating capital quantity and quality (vintage) as distinct, independently evolving state variables. It is precisely this heterogeneity in vintage technology—absent in Wang-Zhang's isomorphic framework—that drives our amplification results. While Wang and Zhang (2025) examine how new technologies are created through R&D investment, our framework explains how existing technologies diffuse across heterogeneous firms through lumpy capital investment. Our empirical evidence confirms that investment age proxies for technology vintage and that this mechanism operates beyond intangible capital accumulation.

8 Concluding Remarks

This paper establishes that heterogeneity in investment-specific productivity—arising from timing of capital-embodied technology adoption—constitutes a quantitatively important amplification mechanism for business cycles.

Using Italian firm-level data, we show that investment age (time since last major capital expenditure) proxies for technology vintage and negatively predicts TFP. Each additional year of investment age reduces productivity by about 0.5%. This relationship survives controls for lagged TFP, intangible capital investment, firm trends, and managers' forward-looking sales expectations. Sectoral analysis using SVAR-identified investment-specific shocks confirms causality: positive shocks increase investment contemporaneously and boost TFP with a one-period lag, consistent with time-to-build.

We develop a general equilibrium model where firms optimally time adoption of successively more efficient investment goods, subject to non-convex switching costs. Unlike standard vintage capital models (Solow, 1960; Greenwood, Hercowitz and Krusell, 1997)

where all firms access the frontier, or lumpy investment models (Khan and Thomas, 2003, 2008) where all firms access identical technology, our framework generates endogenous heterogeneity in technology *quality*. Compositional shifts in the vintage distribution—as laggard firms upgrade during booms—generate productivity improvements that amplify and propagate shocks beyond their initial impulse.

The vintage model replicates key microeconomic patterns: cross-sectional distributions of investment rates and investment age, near-zero autocorrelation of investment rates, and approximately 10% of observed TFP dispersion. For business cycles, the model requires investment-specific shocks roughly 50% smaller than a neoclassical growth benchmark to match observed volatility, while delivering greater persistence. When technology quality varies endogenously with adoption timing, microeconomic heterogeneity materially affects aggregate dynamics—overturning the "irrelevance" finding of Khan and Thomas (2003). Policies facilitating adoption, such as investment tax credits, may have larger effects than standard models suggest, as they shift the vintage distribution rather than merely aggregate capital quantity. The amplification mechanism implies investment-specific shocks contribute more substantially to business cycles than models without vintage heterogeneity estimate.

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.

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

Blanchard, Olivier Jean, and Danny Quah. 1989. "The dynamic effects of aggregate demand and supply disturbances." *American Economic Review*, 79(4): 655–673.

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.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2012. "Americans Do IT Better: US Multinationals and the Productivity Miracle." *American Economic Review*, 102(1): 167–201.

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

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *Quarterly Journal of Economics*, 117(1): 339–376.

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

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

Cloyne, James, Clodomiro Ferreira, Maren Froemel, and Paolo Surico. 2023. "Monetary Policy, Corporate Finance, and Investment." *Journal of the European Economic Association*, 21(6): 2586–2634.

Comin, Diego, and Mark Gertler. 2006. "Medium-Term Business Cycles." *American Economic Review*, 96(3): 523–551.

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.

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, Colleen Lipa, and William Wu. 2025. "Investment as a Source of Productivity Growth." *FEDS Notes*. Washington: Board of Governors of the Federal Reserve System.

Fisher, Jonas D. M. 2006. "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks." *Journal of Political Economy*, 114(3): 413–451.

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

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

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

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

Guerrieri, Luca, Dale Henderson, and Jinill Kim. 2020. "Interpreting shocks to the relative price of investment with a two-sector model." *Journal of Applied Econometrics*, 35(1): 82–98.

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

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

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

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

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

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

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

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

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

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

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

Koby, Yann, and Christian K. Wolf. 2019. "Aggregation in heterogeneous-firm models: A sufficient statistics approach." Princeton University Technical report.

Lev, Baruch, and Theodore Sougiannis. 1996. "The Capitalization, Amortization, and Value-Relevance of R&D." *Journal of Accounting and Economics*, 21(1): 107–138.

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

McElheran, Kristina, J. Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia Foster, and Nikolas Zolas. 2024. "AI adoption in America: Who, what, and where." *Journal of Economics & Management Strategy*, 33(2): 375–415.

Nadiri, M. Ishaq, and I. R. Prucha. 1996. "Estimation of the Depreciation Rate of Physical and R&D Capital in the U.S. Total Manufacturing Sector." *Economic Inquiry*, 34(1): 43–56.

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.

Romer, Paul M. 1990. "Endogenous Technological Change." *Journal of Political Economy*, 98(5, Part 2): S71–S102.

Runkle, David E. 1987. "Vector Autoregressions and Reality." *Journal of Business and Economic Statistics*, 5: 437–442.

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

Sedláček, Petr. 2020. "Creative Destruction and Uncertainty." *Journal of the European Economic Association*, 18(4): 1814–1843.

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

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

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

Wang, Lixing, and Donghai Zhang. 2025. "A Theory of Lumpy Innovation and Aggregate Dynamics." SSRN Working Paper 4818040.

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

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

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

A Data Sources

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

Table *A*.1: Sectoral Data

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

Electricity and gas supply includes industry 35. The water supply sector includes industries 36-39. The construction sector includes industries 41-43. The wholesale and retail trade sector includes industries 45-47. The transportation and storage activities sector includes industries 49-53. The accommodation and food service sector includes industries 55 and 56. The information and communication sector includes industries 58-63. The financial and insurance activities sector includes industry 66. The real estate activities sector

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

B Investment Rates and Total Factor Productivity

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

B.1 Total Factor Productivity

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

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

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

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

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

value-added, as defined by

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

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 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 The Lumpy Nature of Capital Accumulation

We now document the lumpy nature of capital accumulation at the firm level by computing the distribution of investment rates for the Cerved database. 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.

Investment $I_{f,t}$ includes expenditures on equipment and structures as they are not separately identifiable in our data. Table A.2 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 that exhibit a positive spike—i.e., an investment rate above 20%. On average, these investors account for about two thirds of total investment. Firms that experience small capital adjustments (defined as in Øivind and Schiantarelli (2003) as experiencing $ik_{f,t}$ between negative 5 and 5%) account for only 6% of total investment.[18] When grouping firms by their investment rates, the share of investment that the individual groups of firms account for differs substantially; instead, the shares of output and employment are similar.

D Structural VAR Analysis

We estimate a structural vector autoregression (SVAR) to establish two objectives. First, we assess the quantitative importance of investment-specific productivity shocks in driving

[18]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.

Table *A*.2: Cross-Sectional Distribution of Investment Rates

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

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

business cycle fluctuations in the Italian economy. Second, we use the identified structural disturbances to conduct a conditional analysis of how firm- and sectoral-level total factor productivity (TFP) responds to these aggregate investment-specific shocks, thereby linking macro dynamics to micro outcomes.

Following the identification strategy of Fisher (2006) and Guerrieri, Henderson and Kim (2020) for U.S. data, we estimate a five-variable VAR with two lags using quarterly Italian data. The system includes: (i) the growth rate of the relative price of investment (RPI), constructed as the log-differenced ratio of the price deflator for machinery, equipment, weapons systems, and durable consumption goods to the price deflator for household consumption of semi-durable goods, non-durable goods, and services; (ii) labor productivity growth, measured as the log-differenced ratio of real GDP to total hours worked; (iii) hours per capita, defined as the log ratio of total hours worked to the civilian population; (iv) real investment growth per capita, computed as the log-differenced ratio of real gross fixed capital formation to population; and (v) real consumption growth per capita, calculated as the log-differenced ratio of final household consumption expenditure to the civilian population.

The sample period is chosen to balance data availability with structural stability. The 1998:Q1 starting date ensures consistency with availability of data from national accounts aggregates, while the 2014:Q2 endpoint avoids potential structural breaks associated with the European Central Bank's unconventional monetary policies and the effective lower bound on interest rates during and after the sovereign debt crisis. We identify structural shocks using long-run restrictions in the spirit of Blanchard and Quah (1989) and Fisher (2006). Specifically, the identification scheme imposes two key restrictions: first, only a shock to the relative price of investment can move that price permanently; second, only shocks to the relative price of investment and to labor productivity can move the level of labor productivity permanently. All other shocks remain unidentified. The long-run restrictions allow us

to distinguish between investment-specific productivity shocks that permanently affect relative prices and neutral productivity shocks that permanently affect labor productivity but not relative prices. The thick dashed lines in Figure A.1 show the estimated effects of a one-

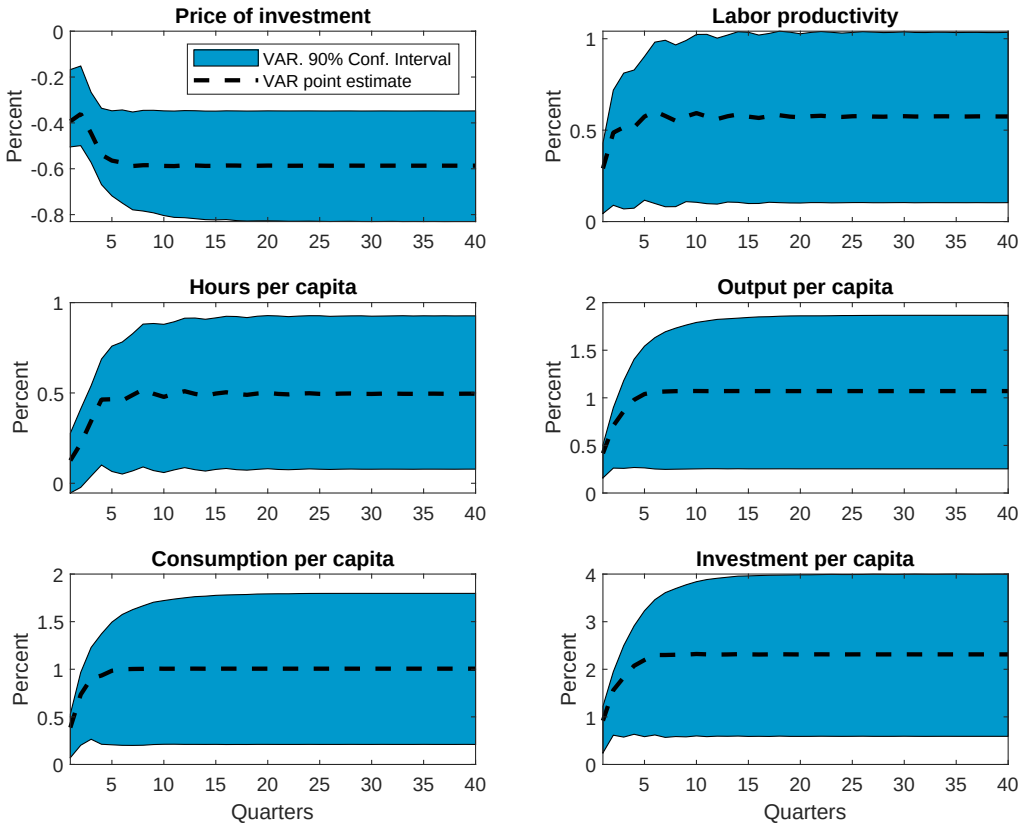


Figure A.1: Impulse response functions following a permanent decrease to the relative price of investment.

standard-deviation shock that permanently reduces the relative price of investment. The point estimate indicates a decline of approximately 0.5 percentage points. The shaded areas denote 90% confidence intervals constructed following Runkle (1987), based on 1,000 bootstrap replications. As is typical for identification schemes based on long-run restrictions, the confidence intervals are wide. Nonetheless, they indicate statistically significant and economically meaningful responses for consumption, investment, and output. Moreover, hours per capita increases significantly with a lag. Taken together, these estimates suggest that investment-specific shocks generate substantial fluctuations in macroeconomic aggregates, consistent with the findings of Guerrieri, Henderson and Kim (2020) for the United States.

Table A.3 documents the contribution of investment-specific shocks to the forecast error variance of investment and consumption across different horizons. Investment-specific

shocks account for about 30% of the variance in these variables. These results corroborate the findings of Guerrieri, Henderson and Kim (2020) for the U.S. economy and underscore the importance of investment-specific technological change as a driver of aggregate fluctuations.

Finally, given that the SVAR is estimated using quarterly data while the sectoral data are at annual frequency, we construct the annual investment-specific productivity shock by summing the quarterly shocks over each calendar year.

Table *A*.3: Historical Variance Decomposition

	Forecast Horizon (Quarters)					
Variable	1	4	8	12	20	40
Panel A: Investment-Specific Productivity Shock						
Investment Per Capita	30.19	33.02	32.51	32.39	32.34	32.33
Consumption Per Capita	26.13	35.13	35.21	35.20	35.20	35.20
Relative Price of Investment	49.04	42.45	41.95	41.81	41.74	41.73
Labor Productivity	11.49	9.71	8.05	7.60	7.41	7.39
Hours Per Capita	2.56	5.12	4.10	3.90	3.82	3.81
Panel B: Neutral Productivity Shock						
Relative Price of Investment	29.30	31.07	31.09	31.05	31.03	31.03
Investment Per Capita	6.39	12.05	12.84	12.86	12.88	12.88
Consumption Per Capita	25.60	24.46	24.39	24.39	24.39	24.39
Labor Productivity	46.56	33.41	30.36	29.70	29.44	29.41
Hours Per Capita	11.54	21.26	19.92	19.96	20.00	20.00

Notes: Entries report the percentage of forecast error variance at various horizons explained by each structural shock. The five-variable VAR includes the growth rate of the relative price of investment, labor productivity growth, hours per capita, real investment per capita growth, and real consumption per capita growth. The sample period is 1998:Q1-2014:Q2. Identification uses long-run restrictions: investment-specific productivity shocks are the only shocks that permanently affect the relative price of investment; both investment-specific and neutral productivity shocks permanently affect labor productivity. The VAR includes four lags.

We also estimate an expanded specification that includes the real interest rate maintining the same identification scheme. This specification complements prior work that uses the real interest rate response to discriminate among propagation mechanisms in the lumpy investment literature. For instance, Winberry (2021) finds that the response of the real interest rate to a *neutral* total factor productivity shock contrasts with standard RBC models and uses this evidence to discriminate across models of lumpy investment. We measure the real interest rate as the three month interest rate on Italian treasury bill minus headline inflation in four-quarter changes. Figure *A*.2 indicates that the estimated response of the real interest rate, although not statistically significant, aligns with what is observed in the model. The real interest rate increases in response to an increase in investment-specific productivity (or, equivalently, a decrease in the relative price of investment) after the first quarter. This result is consistent with the model's predictions at annual frequency, where the real interest rate increases over the response horizon. Importantly, following an investment-specific shock,

the real interest rate peaks at quarter six, reaching 7 basis points, a magnitude comparable to the real interest rate response in the model.

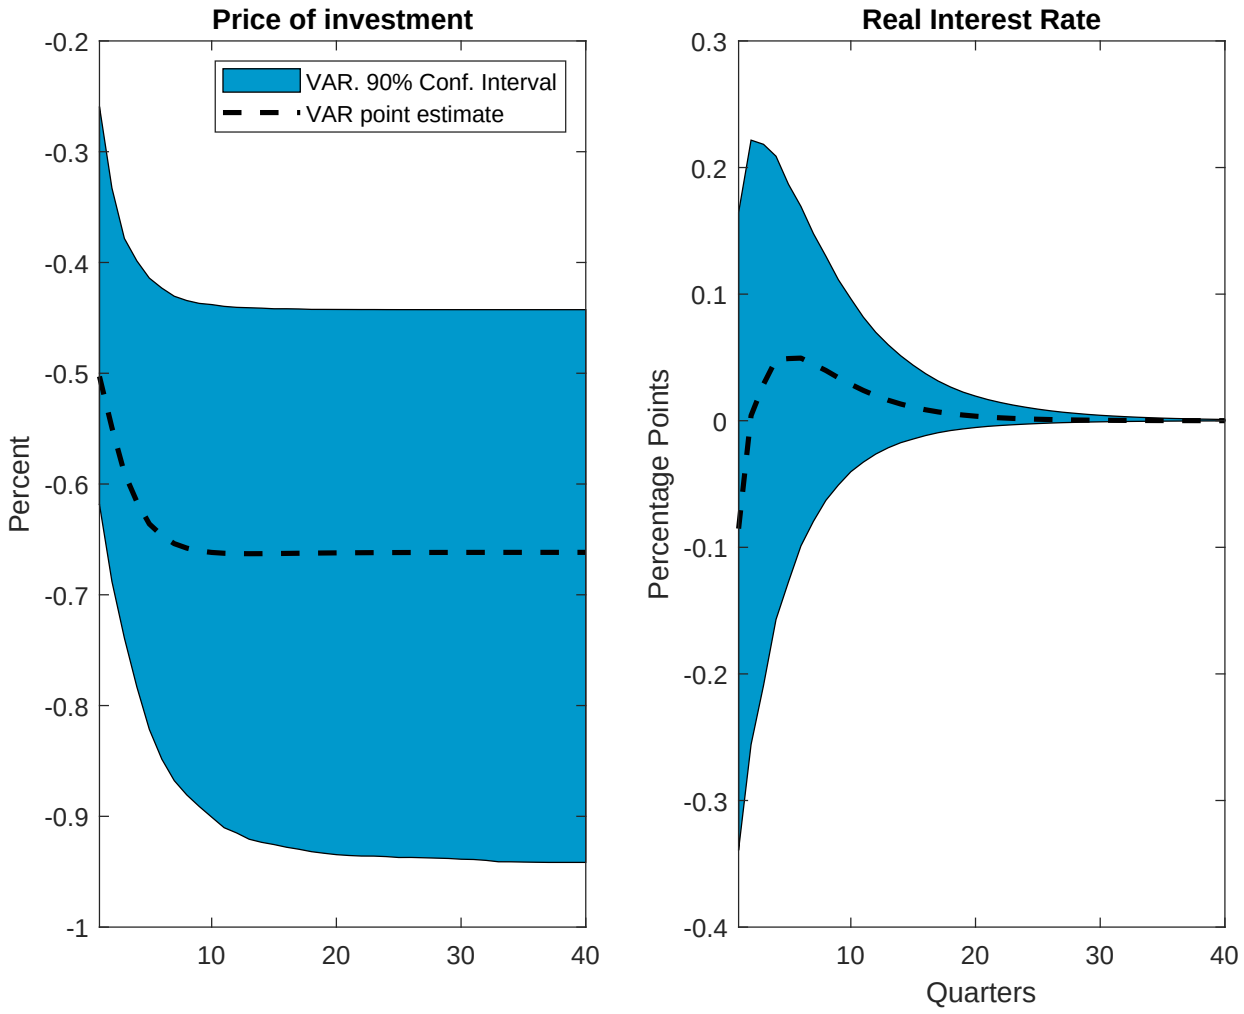


Figure $A.2$: Impulse response of the real interest rate following a permanent decrease in the relative price of investment (VAR with expanded specification).

E Equilibrium and (S,s) Decision Rules

Given the presence of fixed cost, the adoption and investment decisions are akin to exercising an option. Consider a firm of a type $(\varepsilon, z, \tilde{k})$ drawing adjustment cost ξ. Define the value

associated with the value of action V^A and the one with the inaction choice V^{NA} as

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

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

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

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

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

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

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

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

where it is understood that i^A and i^{NA} depend upon the firm's current state. Finally, we turn to the evolution of the firm distribution, $\mu_{t+1} = \Gamma(q_{0,t}, \mu_t)$. It is useful to define the indicator

function $J(x) = 1$ if x is true, and 0 otherwise. For each $(\varepsilon_t, q_{j,t}, \tilde{k}_t) \in S$

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

F Computation

F.1 Solution Algorithm

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

We estimate a regression conditional on each realization of the aggregate process, $q_{0,t}$. In Tables A.4 and A.5, we assess the accuracy of the forecasting rule for both models. For simplicity, we report only one set of results for the neoclassical growth benchmark model as the accuracy is virtually the same.

Table $A.4$: Forecasting Rules - Vintage Model

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

We find that the algorithm yields a very accurate solution as reflected in the high R^2 and small standard errors. As discussed by Den Haan (2010), the values of R^2 are averages and scaled by the variance of the dependent variable, so the lower variance of $\overline{Q}_{t+1}$ relative to the other approximated variables explains its somewhat lower R^2 even for a similar mean and maximum errors. To provide a robust statistic, we perform a long simulation of the model (5,000 periods) and compare the paths of the approximated variables with an alternative simulation obtained iterating only on the estimated law of motions. We compute the maximum and mean distance of each approximated variable from the value taken in the actual simulation. These values are reported in Table A.6. As all variables are in logs, the errors can be interpreted as percentage deviations. The results show that the solution is not only accurate in the sense that it produces accurate one-step ahead forecast, but also in that the forecast errors do not accumulate over time.

Table $A.5$: Forecasting Rules - RBC Model

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

Table $A.6$: Forecasting Rules - Den Haan Test

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

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

G Business Cycle Moments: Partial Equilibrium

We now examine the business cycle moments of our economy under partial equilibrium, comparing the vintage model with its neoclassical counterpart. This exercise demonstrates that general equilibrium forces significantly shape the dynamics of lumpy investment, consistent with the findings of Thomas (2002) and Khan and Thomas (2008).

Table *A.7*: Unfiltered Business Cycle Statistics: Technology Shock

	GDP	Consumption	Investment	TFP	Employment
	(A)	(B)	(C)	(D)	(E)
HP-FILTERED MOMENTS					
Baseline Vintage - General Equilibrium					
σ_X	0.56	0.32	3.60	0.06	0.57
$Corr(X_t, X_{t-1})$	0.89	0.88	0.81	0.79	0.80
$Corr(X_t, GDP_t)$	1	0.24	0.88	0.86	0.84
Baseline Vintage - Partial Equilibrium					
σ_X	2.79	n.a.	54.05	0.09	2.79
$Corr(X_t, X_{t-1})$	0.74	n.a.	-0.09	0.68	0.74
$Corr(X_t, GDP_t)$	1	n.a.	0.17	0.95	1
UNFILTERED MOMENTS					
Baseline Vintage - General Equilibrium					
σ_X	0.82	0.49	4.44	0.24	0.69
$Corr(X_t, X_{t-1})$	0.94	0.94	0.87	0.96	0.86
$Corr(X_t, GDP_t)$	1	0.54	0.94	0.90	0.80
Baseline Vintage - Partial Equilibrium					
σ_X	3.47	n.a.	55.04	0.23	3.45
$Corr(X_t, X_{t-1})$	0.83	n.a.	-0.05	0.93	0.82
$Corr(X_t, GDP_t)$	1	n.a.	0.12	0.96	0.99

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

H Unfiltered Business Cycle Moments

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

Table *A*.8: Unfiltered Business Cycle Statistics: Technology Shock

	GDP	Consumption	Investment	TFP	Employment
	(A)	(B)	(C)	(D)	(E)
Baseline Vintage					
σ_X	0.82	0.49	4.44	0.31	0.69
$Corr(X_t, X_{t-1})$	0.94	0.94	0.87	0.92	0.86
$Corr(X_t, GDP_t)$	1	0.54	0.86	0.90	0.80
Neoclassical model $\overline{\xi}_{l,h} = 0$					
σ_X	0.51	0.36	3.97	0.10	0.61
$Corr(X_t, X_{t-1})$	0.78	0.74	0.57	0.78	0.56
$Corr(X_t, GDP_t)$	1	0.06	0.84	0.66	0.80

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

Appendix References

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

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.

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.

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

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

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

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

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

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

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