Production Networks and Stock Returns: The Role of Creative Destruction*

Michael Gofman*, Gill Segal‡, and Youchang Wu§

July 31, 2017

Abstract

We establish a novel return spread based on the distance between firms and final consumers in a production network. Firms with the longest distance to consumers earn an excess monthly return of 105 basis points relative to final goods producers. We explain this spread quantitatively using a general equilibrium model with multiple layers of production. The driving force behind the spread is creative destruction, which reduces firms' exposure to productivity shocks. The spread is smaller for firms that belong to supply chains with lower competition. Overall, our results demonstrate a novel effect of creative destruction on firms' cost of capital.

Keywords: production networks, stock returns, creative destruction, monopolistic competition, technological innovations

*This paper benefited from comments and suggestions by Enghin Atalay, Mikhail Chernov, Riccardo Colacito, Max Croce, Dean Corbae, Phil Dybvig, Roberto Gutierrez, Brandon Julio, Mete Kilic, Robert Novy-Marx, Alexi Savov, Ivan Shaliastovich, Amir Yaron and seminar participants at UW-Madison, University of Oregon and Goethe University Frankfurt. We are grateful to Pavel Brender and Alexander Dentler for excellent research assistance. Michael Gofman acknowledges financial support from Wisconsin Alumni Research Foundation (WARF). All errors are our own.

*michael.gofman@simon.rochester.edu, University of Rochester.

‡Gill_Segal@kenan-flagler.unc.edu, University of North Carolina at Chapel Hill.

§ywu2@uoregon.edu, University of Oregon.

1 Introduction

A growing literature in macroeconomics and finance examines the joint dynamics of firm-level production-based characteristics and stock returns. It has been shown that investment rate, hiring growth, and productivity correlate with firms' expected cost of capital.[1] To date, most of the research that connects investment behavior and stock returns tends to treat all firms as producing the same homogeneous goods. A few studies explicitly model heterogeneous goods and multi-sector economies and derive implications for risk premia (see, e.g., Gomes et al. (2009), and Kogan and Papanikolaou (2014)). These latter papers use crude granularity by dividing firms into producers of durables versus non-durables, or consumption goods versus capital goods. However, in reality, the majority of firms produce differentiated intermediate inputs used by firms in long production chains. These production chains constitute a complex production network that transforms raw materials into final products via multiple layers of production.[2] How risk is distributed along production chains remains an open question.

Intuitively, not all firms along the same production chain should benefit equally from common technological advancements. Innovation which benefits the production of new vintages of capital, produced by suppliers at the top of the production chain, can devalue old capital of customer firms located at the bottom of the chain. This creative destruction yields differential exposure of firms along the multi-chain production network to aggregate shocks. Motivated by this intuition, we seek to empirically and theoretically study the relationship between the *granular* location of a firm in a multi-chain production network and its stock returns.

The main empirical challenge to study this relationship is the availability of comprehensive data that allows one to measure a firm's position in the production network. A common practice in the literature is to study production networks using industry level input-output tables, ignoring any intra-industry variation across firms. We overcome this challenge by exploiting a novel database of supplier-customer relationships. The database allows us to compute firm's upstreamness in a production network at a monthly frequency. We use this measure to compare stock returns of firms at different vertical positions in a production network.

To compute a firm's upstreamness measure, we decompose a production network into

[1] See discussion in the related literature section below.

[2] A famous representation of this multi-layer structure of production is the so-called "Hayekian Triangle" (Hayek (1935), page 39).

layers of production. All firms in the same layer are separated by the same number of supplier-customer links from the bottom layer firms that produce final consumption goods. We define a vertical position of any firm as the smallest number of supplier-customer links between itself and firms at the bottom layer. A firm's vertical position is the primary production-based characteristic of interest in this paper.

The first contribution of this paper is to empirically document a novel, strong, and monotonic link between a firm's distance to consumers (i.e., vertical position) and its cost of capital. We show that expected returns are higher for firms that are further away from final goods producers. An investment strategy that longs firms with the longest distance to consumers (highest vertical position) and shorts firms with the shortest distance (lowest vertical position) earns a return of 105 basis points per month. We refer to this difference in returns as the TMB (top-minus-bottom) spread. This spread remains significant after we control for the common risk factors. It is also orthogonal to existing cross-industry spreads such as Durables-minus-Non-Durables (see Gomes et al. (2009)), or Investment-minus-Consumption (see Kogan and Papanikolaou (2014)).

Our second contribution is to build a production-based asset-pricing model, which demonstrates how the force of creative destruction affects asset prices and gives rise to the TMB spread. The model provides a quantitative risk-based explanation for the spread. It also delivers a novel insight that contrary to common wisdom, creative destruction may reduce firms' cost of capital. The model ingredients include multiple production layers (a production chain), an aggregate productivity shock that features time-varying and persistent growth rate, and a household with Epstein and Zin (1989) and Weil (1989) preferences. In the model, the output of a representative firm in the top layer, layer N, is used to supply intermediate capital to the layer below it, layer $N-1$, which in turn uses this capital stock (and labor) to produce and supply intermediate capital goods to a representative firm in layer $N-2$. The production chain uses intermediate capital goods from one layer as inputs to the next layer until a representative firm at the bottom layer uses its capital stock and labor to produce final consumption goods. The markets for labor and capital are competitive.

We calibrate the model to target macroeconomic moments. The calibrated model exhibits a monotonic relationship between stock returns and vertical position and is able to generate a 12% per annum spread between the top and bottom layers, which quantitatively matches its empirical counterpart. Moreover, in both the model and the data the largest return differential is between the top two layers. The model also successfully

matches moments related to the market portfolio and the risk-free rate. The sizable return spread between the top and bottom layers are due to differential exposures to aggregate productivity shocks, which decline monotonically from the top to the bottom layer.

The pattern of exposure to aggregate productivity across the layers results from the force of creative destruction in the spirit of Schumpeter (1942). Intuitively, a positive productivity shock has a dual effect on firm valuation. On one hand, it acts as a positive demand shock for the firm's output, which implies higher future cash flows or improved growth options.[3] On the other hand, a positive productivity shock also acts as a positive supply shock for the firm's intermediate capital input. This supply shock lowers the value of the firm's assets-in-place. Put differently, technological improvements make the production of firms' capital input easier and cheaper, which erodes the marginal value of their existing stock of capital.

A firm at the bottom of the production chain experiences the greatest impact of the latter effect, because its existing capital is (effectively) built using the capital goods produced by all the layers above it, which encompass the entire production chain. As each of the intermediate capital goods becomes cheaper to produce, the value of the assets-in-place of the bottom layer firm drops the most, as its replacement cost of capital falls drastically. By contrast, the firm at the top of the production chain has no suppliers. As such, it is not subject to this creative destruction force. Creative destruction makes the sensitivity of a firm to productivity shocks less positive, acting as a hedge. Since top layer firms do not experience the Schumpeterian effect, their total exposure to aggregate productivity is larger than that of bottom layer firms. Firms in the middle layers experience some amount of devaluation to their installed capital following a positive productivity shock, but it is not as large as that of the bottom layer firms because assets of the firms in the middle layers are produced using inputs from fewer layers. This explains not only the positive spread between the top and the bottom layers (TMB spread), but also why expected returns monotonically increase with the vertical position.

We further test whether the spread indeed reflects the force of creative destruction by augmenting our model to accommodate monopolistic competition. The augmented model provides additional predictions that allow us to test the economic mechanism. The model predicts that the TMB spread is smaller when firms at each layer of production have greater monopolistic power. The monopolistic power enhances firms' growth options, which makes

[3]By 'growth options' in this paper, we refer to any component of the firm's value that is not related to the valuation of its installed capital, or assets-in-place. This includes the firm's future monopolistic rents, if exist.

them more exposed to productivity shocks. In addition, assets-in-place of the bottom layer firms devalue less when there is less competition because monopolistic suppliers keep part of the benefits from the technological improvements by not reducing the prices of their output as much as competitive suppliers would. With a diminished hedging of creative destruction, firms at the bottom layer become more exposed to the shocks, their expected returns increase, decreasing the TMB spread.

To test predictions of the augmented model, we compute a novel measure of supply chain competition, which takes into account the number of competitors at each level of production. This measure combines the network framework of supplier-customer relationships with information about the number of competitors each firm has. A firm has a high measure of "supply chain competition" if it has many competitors, if its suppliers have many competitors, if the suppliers of its suppliers have many competitors, and so on.

We split the sample into two subsamples based on the measure of a firm's supply chain competition relative to that of the median firm at the same vertical position. Consistent with the model's prediction, we find that the spread is smaller in the sample of firms that face a low level of supply chain competition. Moreover, we show in the subsample of high supply chain competition that the TMB spread has a positive and statistically significant beta with respect to quarterly changes in the utilization-adjusted total factor productivity growth (TFP). In the low competition subsample, the beta of the spread to TFP shocks is indistinguishable from zero. These results further confirms the prediction of our model. We conclude that supply chain competition plays an important role in determining the magnitude of the spread, providing further confirmation that the spread is driven by the force of creative destruction.

Related theoretical literature. Technological advancements are an important driver of economic growth. However, not all firms benefit equally from innovation, a notion that traces back to Schumpeter (1942). Gârleanu et al. (2012), Loualiche (2016), and Barrot et al. (2016) study displacement risk, which refers to the notion that innovation can benefit new firms or entrants at the expense of incumbent firms. In our model, common technology improvements benefit all firms, but downstream firms suffer more from the depreciation of their installed capital due to creative destruction. As a result, downstream firms are less exposed to the shock, which explains why the TMB spread within incumbent firms is positive and significant.

We contribute to the extensive literature about creative destruction started by Schumpeter (1942). From the theoretical modeling perspective, creative destruction has been used to

study endogenous growth (Aghion and Howitt, 1992), effects of labor market frictions on the resource allocation process (Caballero and Hammour, 1996), and even to explain "animal spirits" (Kennedy and Lloyd-Ellis, 2003). The novelty of our model is to explicitly account for the multi-layer nature of production process, rather than to assume a representative firm. Our model reveals a counterintuitive result that creative destruction can lower firms' cost of capital by making them less exposed to productivity shocks.

Our paper is also related to the literature that examines cross-sector return spreads. Papanikolaou (2011), Kogan and Papanikolaou (2014), Garlappi and Song (2016), and Yang (2013) examine the cross-sectional pricing difference between the consumption sector and the investment sector. We deviate from this literature by using a network-based methodology to compute individual firms' vertical positions. Quantitatively, the majority of the TMB spread stems from *within* the investment sector, not from the return differentials of consumption versus investment firms. Gomes et al. (2009) document higher expected returns for firms that produce durables than for non-durables producers. The TMB spread is orthogonal to this cross-industry spread.

More broadly, our paper contributes to the theoretical literature which studies macroeconomic dynamics jointly with aggregate asset-pricing implications, or cross-sectional return risk premia, including Jermann (2010), Berk et al. (1999), Tallarini (2000), Boldrin et al. (2001), Gomes et al. (2003), Carlson et al. (2004), Zhang (2005), Belo et al. (2014), among others. Our contribution is to demonstrate how creative destruction generates differential exposures to productivity shocks, yielding a quantitatively large spread.

Our paper is also related to the macro literature that studies the effect of the structure of production networks on aggregate fluctuations (Acemoglu et al., 2012; Bigio and La'O, 2016; Acemoglu et al., 2017; Atalay, 2017). While this literature investigates how idiosyncratic shocks aggregate in the network, our paper investigates how firms differentially exposed to aggregate shocks because of the network effects.

Related empirical literature. The paper contributes to the empirical literature about the effects of creative destruction.[4] Very few papers have studied the implications of creative destruction on stock returns. Hobijn and Jovanovic (2001) argue that improvements in information technologies (IT) are responsible for the the stock market decline in early 1970s. Chun et al. (2008) argue that improvements in the IT sector are responsible for the elevated firm performance heterogeneity in the late 20th century. In our paper, we do not try to explain fluctuations in the aggregate stock returns or a cross-sectional

[4]See Caballero (2008) for an excellent survey.

heterogeneity in firm performance due to innovations in a particular sector, but rather to explain why and to what extend some firms are more affected by aggregate productivity shocks than others. To the best of our knowledge, our paper is the first to quantitatively measure how differential degree of creative destruction affects firms' cost of capital.

Our paper is also closely related to the recent asset pricing literature that utilizes supplier-customer relationships or input-output tables. Cohen and Frazzini (2008) and Menzly and Ozbas (2010) study predictability of stock returns via supplier-customer links. In contrast, we are interested in contemporaneous returns across layers of production. Ahern (2013) finds that industries with a higher network centrality measure have higher returns. We verify that the TMB spread is not explained by the differences in firms' network centrality. Ozdagli and Weber (2017) find sizable network effects in the propagation of monetary shocks. Our paper suggests that because of creative destruction, productivity shocks also have indirect (network) effects on firms' valuations. In a recent paper, Herskovic (2017) derives two risk factors based on the changes in network concentration and network sparsity. In our model, changes in the network structure are not the origin of risk, but they can affect the size of each firm's indirect exposure to a general productivity shock because of the forces of creative destruction.[5]

The rest of the paper is organized as follows. In Section 2 we present the data and our measure of the vertical position. Section 3 includes the main empirical results. The baseline theoretical model with perfect competition is presented in Section 4. In Section 5, we introduce monopolistic competition to test the Schumpeterian nature of the spread. Section 6 discusses robustness checks. We conclude in Section 7.

2 Empirical Measures of Vertical Position

2.1 Data

The data used in our empirical analysis are obtained from several sources. We use CRSP for stock prices, Compustat for accounting data, and the FactSet Revere relationships database for information about suppliers, customers, and competitors.

The FactSet Revere relationships database is the most comprehensive database currently

[5]Other asset pricing implications of production networks were studies by Buraschi and Porchia (2012); Aobdia et al. (2014); Branger et al. (2017); Rapach et al. (2015), and Richmond (2015). None of these studies has examined the effect of creative destruction on stock returns or the relationship between stock returns and firms' vertical position.

available that covers a large cross-section of public and private firms. Our sample period is from April 2003, when the database started, to September 2013, when we purchased it from FactSet Revere. The database includes only public information reported by firms and their trading partners or competitors. Regulation SFAS No. 131 requires firms to report customers with more than 10 percent of sales. Some companies voluntarily report additional customers that are below that threshold.[6] FactSet complements this information with additional sources, making the database more comprehensive than the commonly used Compustat's segment data. These sources include SEC 10-K annual filings, investor presentations, corporate action announcements, and press releases. Another advantage of using the FactSet Revere relationships data is that it includes not only information about customers, but also about suppliers and competitors. FactSet's analysts analyze each relationship in depth at annual frequency, but the database is updated daily as new information becomes publicly available. To allow for a sufficient time to Factset's analysts to verify and update the supplier-customer relationships, we use only relationships that were present as of December 2012.

2.2 Supplier-Customer Relationships

We construct a dynamic production network using information about individual links between a supplier and a customer. The unit of observation in the database is a relationship between two firms. We observe a relationship's start and end dates. To get the most comprehensive information about the production network, we combine both the information about customers disclosed by firms and information about their suppliers. A relationship does not need to be reported by both firms to be recorded. For example, if Mellanox Technologies Ltd discloses IBM as a customer, we use this relationship even if IBM has not reported Mellanox Technologies Ltd as its supplier.

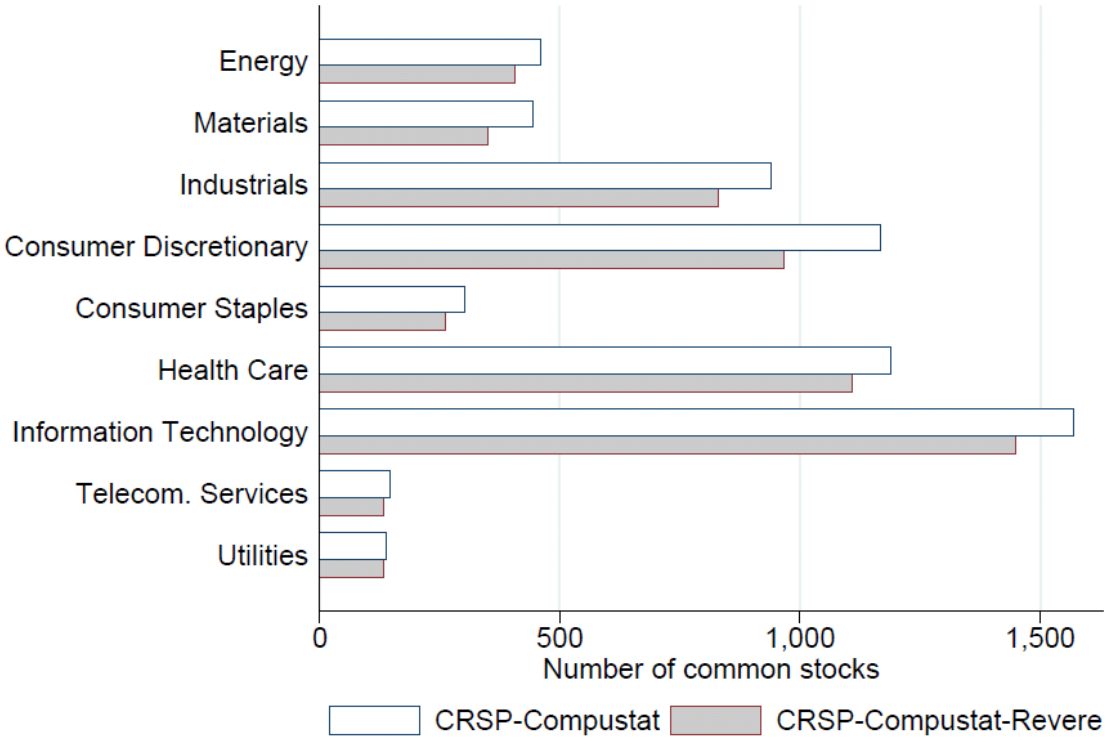
The database includes 433,271 supplier-customer relationships between 193,851 pairs of firms, covering a total of 43,656 firms. These relationships require cleaning. Our cleaning procedure includes removing duplicate records, removing redundant relationships whose start and end dates are within the time period of a longer relationship between the same pair of firms, and eliminating the gap between two relationships between a single pair of firms when the second starts within six months of the end of the first. This data cleaning results in 206,264 supplier-customer relationships.

[6]While for some suppliers we observe the percentage of revenues from a given customer, we do not utilize this data in our analysis because it only covers a small subset of supplier-customer relationships.

We merge this sample of firms with firms in the Compustat North America database using CUSIP codes, and are able to find 9,117 matches. We exclude financial firms (GICS code: 40) and industrial conglomerates (GICS: 201050), and end up with 7,801 Compustat firms with at least one supplier-customer relationship. We further match this sample of firms to the CRSP monthly stock database using the CRSP-Compustat linking table. After excluding penny stocks (stocks with a price of less than \$1 in the previous month), our final Revere-Compustat-CRSP matched sample consists of 5,645 common stocks (with CRSP share codes 10, 11, and 12).

Over the 2003-2013 period, the total number of non-penny, non-financial, and non-conglomerate common stocks in the CRSP-Compustat merged database is 6,437. Therefore, our final sample includes about 88% of stocks in this most commonly used stock database. Figure 1 shows the number of stocks in our Revere-CRSP-Compustat matched database vs. the number of stocks in the CRSP-Compustat merged database for each industry. It demonstrates that all industries are well-represented in our sample.

Figure 1: **Sample Coverage by Industry**



The figure presents the number of common stocks in our Revere-CRSP-Compustat matched database vs. the number of common stocks in the CRSP-Compustat merged database for each industry. The industries are classified based on the Global Industry Classification Standard (GICS).

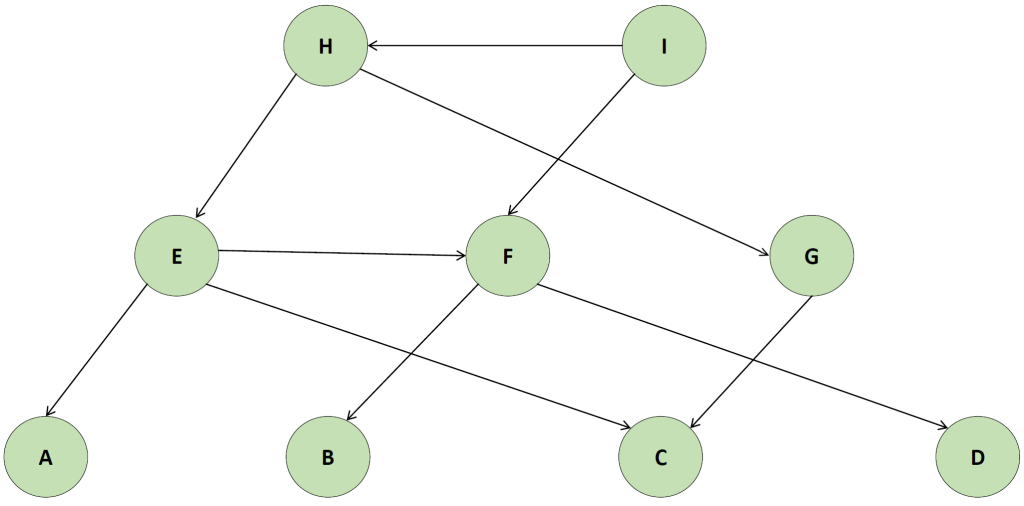
2.3 The Vertical Position Measure

We now describe the main production-based characteristic of interest, a firm's vertical position. The measure relates to the position of a given firm in a production chain (that is, its upstreamness). In this section we illustrate and define this measure, and explain how we compute it from the data.

Production networks can be split into tranches of firms, with firms in the same tranche having a similar distance from consumers (or equivalently, from final consumption good producers). We refer to these tranches as "production layers." The firms at the bottom layer of a network produce final consumption goods. All other firms are direct or indirect suppliers to bottom layer firms. We define the vertical position of any firm as the smallest number of supplier-customer links between itself and firms at the bottom layer.

To illustrate our approach, consider the network of firms depicted in Figure 2. Assume that firms A, B, C and D produce final consumption goods. These firms are, by our convention, operating at layer zero (the bottom layer). The other firms are connected to the bottom layer by one or more supplier-customer links. We denote the (minimal) number of such links as a firm's "vertical position." For instance, Firms E, F and G belong to the same layer, as they have a vertical position of one. Firms H and I, which operate at the top layer, have a vertical position of two.

Figure 2: **Production Network: an Illustration**



The vertical position of each firm is determined endogenously relative to layer zero. Firms in the first layer supply to at least one firm in layer zero. Following the same intuition, firms in layer i supply to at least one firm in layer $i-1$ and to none of the firms

in layers zero to $i-2$. The number of layers of production in each month depends on the supplier-customer relationships and the firms in layer zero.

Formally, consider a distance matrix D_t with n_t rows and m_t columns, where n_t is the total number of firms in the production network in month t and m_t is the number of final goods producers in month t. An element $D_t(i,j)$ of this matrix measures the number of supplier-customer links between firm i and final goods producer j. Given this distance matrix D_t, the vertical position is defined as the minimum number of supplier-customer relationships to any final goods producer:

$$VP_{i,t} = \min_{j \in \{k: VP_{k,t}=0\}} D_t(i,j). \tag{1}$$

The vertical position measure is a global measure that depends on the entire network structure. A firm's vertical position can change even if its set of direct suppliers and customers does not.

A firm is assigned a vertical position i because it supplies one or more firms at vertical position $i-1$. For firms at vertical position i, we do not observe a supplier-customer link to any firms in layers 0 to $i-2$. Given that there are thousands of potential links that we could observe, the lack of a single link provides a strong evidence that the firm's location in the production chain is above $i-1$.

Our goal is to compute the expected returns of firms that belong to different layers of production. Given that the production network is dynamic, the vertical position of firms can change over time. Therefore, we compute our vertical position measures at a monthly frequency by utilizing existing supplier-customer relationships that lasted for at least six months before the measure is computed. Our methodology of computing vertical positions is similar to Gofman (2013), extended to a panel of firms.

Specifically, we use the Revere-Compustat matched sample of 9,117 non-financial, non-conglomerate firms to compute firms' vertical positions.[7] We assign a vertical position of zero to all firms in the Consumer Discretionary (GICS code: 25) and Consumer Staples sectors (GICS code: 30). Firms in layer zero belong to the following industries: Automobiles & Components, Consumer Durables & Apparel, Consumer Services, Media, Retailing, Food & Staples Retailing, Food, Beverages & Tobacco, and Household & Personal Products. We then use equation (1) to estimate vertical positions of the remaining firms in the sample.

[7] We exclude conglomerates because their vertical position in the production network is not precisely measured, and keeping them could introduce a bias into other firms' vertical position measures.

3 Vertical Positions and Stock Returns

3.1 Portfolio Formation

We form portfolios by sorting firms according to their vertical positions. When sorting firms into portfolios at the end of month t, we utilize the vertical position computed at the end of month $t-1$. We do so to ensure that public information about supplier-customer relationships are known to investors. We skip the first six months of the FactSet Revere sample to make sure that the vertical position is based on supplier-customer relationships lasting for at least six months.

The number of production layers, and the distribution of firms across these layers is endogenous. In particular, firms need not be allocated equally across the different layers. In fact, as we illustrate below, top layers of production (layers with a high vertical position) may only include very few firms. To reduce the amount of noise due to the small number of firms at the top layers, we assign firms belonging to layers zero to four into a separate portfolio, but we combine all firms with a vertical position five or above into a single portfolio. In all, we obtain six portfolios.

Table 1: **Transition Matrix of the Vertical Position**

Layer in month t	Layer in month $t+1$				
	layer 1	layer 2	layer 3	layer 4	layer 5
layer 1	0.98	0.01	0.00	0.00	0.00
layer 2	0.01	0.97	0.01	0.00	0.00
layer 3	0.01	0.05	0.92	0.02	0.00
layer 4	0.00	0.03	0.06	0.88	0.02
layer 5	0.01	0.04	0.06	0.07	0.83

This table presents the transition probability of a firm's vertical position from one month to another. The matrix of transition probabilities is computed using monthly data from September 2003 to December 2012. Layer zero is not a part of the transition matrix because it does not change given that it is based on the time-invariant Global Industry Classification Standard (GICS) code reported in Compustat.

The vertical position measure is computed at a monthly frequency and firms can move across the layers of production. Table 1 reports the transition probabilities of firms moving across the layers during the sample period. The matrix suggests that vertical positions are rather stable, especially at the lower level. Firms with a vertical position of one (layer 1) have a 98 percent probability of remaining in layer 1 in the next month. The probability is 83 percent for the firms in the top layer. This result means that vertical positions could

potentially be associated with long-term risk profile of firms. However, some transitions do occur, especially across adjacent layers, suggesting that tracking the network structure dynamically is important.

Table 2 reports summary statistics for each layer. It shows a pyramidal shape of production chains: the number of firms decreases almost monotonically from the bottom layer to the top.[8] This is consistent with the shape of the Hayekian Triangle used by Hayek (1935) to depict the structure of multi-layer production.[9] There is no significant difference between the top and bottom layer in terms of size and leverage. Top firms have lower return on assets (ROA) than bottom firms, but there is no difference in ROA in layers above one. We also find that top layer firms hoard more cash, but the relationship between cash holdings and vertical positions is non-monotonic. We find a monotonic relationship between vertical positions and stock returns, the non-monotonic cash holdings would not be able to explain it. Firms at top layers also exhibit higher asset growth. Since, all else equal, firms with lower asset, investment or hiring growth earn a larger risk premium (see, e.g., Belo et al. (2014)), our novel spread is largely orthogonal to growth-related spreads.

Table 2: **Summary statistics by layer**

	N	Assets	ROA	Debt/Asset	Cash/Asset	Asset Growth
layer 5	24	679	0.094	0.194	0.137	0.061
layer 4	74	530	0.094	0.173	0.135	0.046
layer 3	252	524	0.094	0.182	0.149	0.048
layer 2	908	490	0.094	0.147	0.176	0.034
layer 1	694	598	0.098	0.117	0.187	0.024
layer 0	1067	598	0.119	0.219	0.087	0.016
Layer (5-0)		81.41	-0.025***	-0.024	0.050**	0.044***
		(0.97)	(-5.00)	(-1.60)	(2.83)	(7.06)

This table presents summary statistics for each layer. N is the average number of firms in each layer from September 2003 to December 2012. For all other variables, we first calculate the cross-sectional median in a given month, and then report the time series mean. *Assets* is the total book assets in millions of 2009 USD; *ROA* is operating income before depreciation divided by total book assets. *Debt/Asset*, and *Cash/Asset* are the ratios of total debt, cash and cash equivalents to total book assets, respectively. Newey-West t-statistics are reported in the parenthesis.

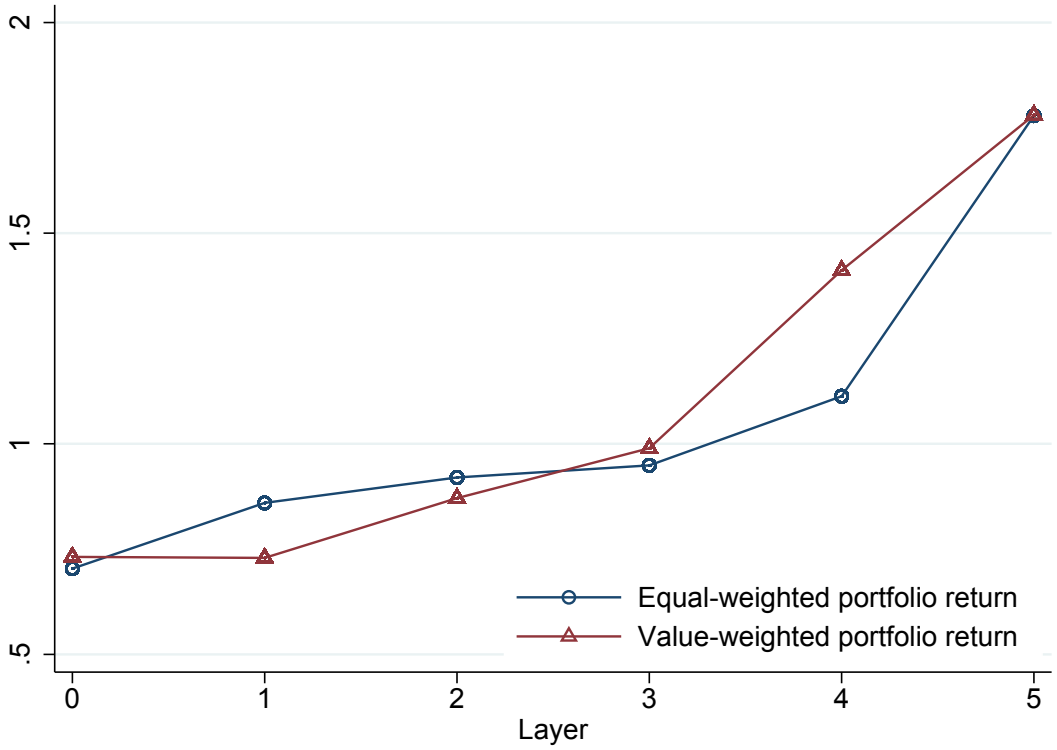
[8]Table 13 in the Online Appendix provides a list of firms that are part of the top layer for at least 18 months.

[9]Our portfolios have different number of stocks because they are based on the production layers. The top portfolio has on average 24 firms and an average total market capitalization of 62 billion USD, which is significantly smaller than the number of firms or a market capitalization of the bottom layer. In Section 6, we verify that the TMB spread is positive and significant even if we combine the top two layers and form five portfolios instead of six. This robustness test alleviates the concern that there is a relatively small number of firms in the top portfolio.

3.2 Portfolio Returns

Figure 3 shows the returns of the value-weighted and equal-weighted portfolios formed based on the firms' vertical positions. The main finding is that expected returns increase in the vertical position. Firms that produce final consumer goods have a nominal value-weighted (equal-weighted) monthly return of 0.73 (0.70) percent, while firms with a vertical position of five or higher have a value-weighted (equal-weighted) average monthly return of 1.78 (1.78) percent.

Figure 3: **Expected Monthly Returns by Vertical Position**



The figure presents the average monthly returns of the value-weighted (red triangles) and equal-weighted (blue circles) portfolios constructed based on firms' vertical positions. Firms with a vertical position of zero are producers of final goods. The sample period is from November 2003 to February 2013.

Table 3 shows the average returns, along with their respective standard deviations, for each layer. Consistently with Figure 3, the returns increase monotonically from the bottom layer to the top layer. The spread between the top and bottom layers is 105 basis points

per month for the value-weighted portfolios and 108 basis points per month for the equal weighted portfolios. Both are economically and statistically significant. The annualized Sharpe ratio of the value-weighted (equal-weighted) TMB portfolio is 0.68 (0.82) from November 2003 to February 2013. During the same period, the Sharpe ratio is 0.39 for the market portfolio, 0.28 for the SMB factor, and 0.29 for the HML factor. The momentum portfolio had a negative average return and a Sharpe ratio of -0.04. The risk-return trade-off of the TMB portfolio is considerably higher than that of the market and more than double the Sharpe ratios of the HML and SMB.

Table 3: **Vertical Position and Stock Returns**

	Value-weighted return		Equal-weighted returns	
	Mean	SD	Mean	SD
layer 5	1.78	6.542	1.779	7.301
layer 4	1.412	6.229	1.113	7.108
layer 3	0.989	5.639	0.949	6.274
layer 2	0.871	4.927	0.92	6.312
layer 1	0.729	4.47	0.86	6.363
layer 0	0.731	3.969	0.704	6.559
spread (5-0)	1.049**	5.359	1.075**	4.535
	(2.07)		(2.51)	

This table presents the summary statistics of the monthly raw returns for each layer and the spread between layers 5 and 0 (the TMB spread). The returns are computed from November 2003 to February 2013.

In Table 4, we test whether the high returns of the TMB portfolio, i.e., a portfolio that longs the top layer and shorts the bottom layer portfolio, can be explained by existing asset pricing models.[10] We consider six well-known factor models, from the classic CAPM, the Fama and French (1993) three-factor model, the Carhart (1997) four-factor model, to the recently proposed Fama and French (2015) five-factor model, the Hou et al. (2015) q-factor model, and the Kogan and Papanikolaou (2014) two-factor model.

The alpha of the spread is not affected by controlling for the factors, and it is still positive and statistically significant at least at the 10% significance level. The TMB portfolio has positive betas with respect to the market excess return, and to the HML, SMB, and MOM

[10]We construct the IMC factor following the procedure in Kogan and Papanikolaou (2014). The consumption and investment industries are identified based on the industry classification developed by Gomes et al. (2009). We thank Ken French, Chen Xue, and Motohiro Yogo for making the factor data and sector classification available on their websites.

Table 4: **Time series regressions of the TMB portfolio return**

	(1)	(2)	(3)	(4)	(5)	(6)
RmRf	0.078	0.016	0.097	0.032	0.007	-0.280*
	(0.68)	(0.11)	(0.69)	(0.21)	(0.05)	(-1.98)
SMB		0.261	0.211	0.353	0.348	
		(1.02)	(0.84)	(1.36)	(1.37)	
HML		0.008	0.119	0.168		
		(0.03)	(0.52)	(0.68)		
MOM			0.227**			
			(2.04)			
RMW				0.292		
				(0.74)		
CMA				-0.621		
				(-1.48)		
I/A					-0.917**	
					(-2.56)	
ROE					-0.069	
					(-0.25)	
IMC						0.784***
						(3.94)
Constant	1.010*	0.993*	0.952*	0.905*	1.100**	1.167**
	(1.98)	(1.93)	(1.88)	(1.66)	(2.15)	(2.42)
R^2	0.004	0.014	0.051	0.045	0.072	0.128

This table presents the results of the time series regressions of the value-weighted Top-Minus-Bottom (TMB) portfolio returns on various risk factors. The TMB portfolio is constructed by taking a long position in the value-weighted portfolio of companies with a vertical position of five or above, and shorting a value-weighted portfolio of companies with a vertical position of zero. The returns are computed from November 2003 to February 2013. RmRf is an excess return of the market portfolio. SMB and HMB are the size (small-minus-big) and book-to-market (high-minus-low) factors in the Fama and French (1993) three-factor model, and MOM is the momentum factor in the Carhart (1997) four-factor model. RMW and CMA are the robust-minus-weak profitability and conservative-minus-aggressive investment factors, respectively, in the Fama and French (2015) five-factor model. I/A and ROE are the investment and profitability factors, respectively, in the q-factor model of Hou et al. (2015). IMC is the investment-minus-consumption factor in the Kogan and Papanikolaou (2014) two-factor model. T-statistics are reported in parentheses. Significance at the 5 and 10 percent levels are indicated by ** and * respectively.

factors, but only the loading on MOM is statistically significant. It has a significantly negative loading on the investment factor of the q-factor model (I/A), and an insignificantly negative loading on the investment factor in the Fama and French (2015) five-factor model (CMA). Due to the negative loading on the investment factor, the alpha of the TMB portfolio in the q-factor model (110 bps per month) is even higher than the average raw

return (105 bps per month). The loading on the investment-minus-consumption (IMC) factor of Kogan and Papanikolaou (2014) is significantly positive, which is to be expected because the bottom layer consists of consumption good producers. However, the alpha of the TMB portfolio is still economically and statistically significant, suggesting that the TMB spread is distinct from the IMC factor. We stress that the TMB spread is mostly within the investment sector, and not a cross-industry spread as IMC.

While not reported, the results for the equal-weighted TMB portfolio spread are qualitatively similar, but quantitatively even stronger. The Carhart (1997) four-factor alpha is 116 basis points per month, while the Fama and French (2015) five-factor alphas is 111 basis points, both statistically significant at the 1% level. These results provide further evidence that the standard factor models are not able to price the cross-section of stock returns sorted on vertical position.

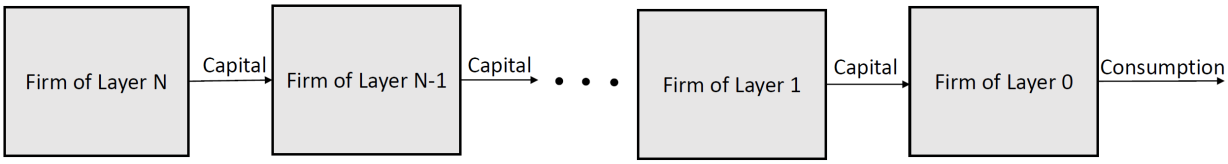
In the next section, we develop a general equilibrium asset-pricing model that is able to shed light on why firms at different vertical positions have different expected returns.

4 General Equilibrium Asset-Pricing Model with Multiple Layers of Production

4.1 The Model

This section describes the general-equilibrium model used to rationalize the production-layer return spread.

There are $N+1$ layers of production in the economy, indexed by $j \in \{0, 1, ..., N\}$. Each production layer is captured by a single representative firm. The firms that operate in layers $\{1, .., N\}$ produce differentiated (intermediate) capital goods. A firm that operates in layer $j \in \{1, .., N\}$ supplies capital to the firm operating in the layer vertically below it, $j-1$. The firm in the bottom layer ($j = 0$) produces final consumption goods, sold to the household for consumption. The economy's production structure is schematically depicted below:



4.1.1 Aggregate Productivity

Aggregate productivity is denoted by Z_t, and its lower case denotes log-units. The log-growth of aggregate productivity features a persistent component as in Croce (2014):

$$\Delta z_{t+1} = \mu_z + x_t + \sigma_z \varepsilon_{z,t+1}, \tag{2}$$

$$x_{t+1} = \rho_x x_t + \phi_x \sigma_z \varepsilon_{x,t+1}, \tag{3}$$

where $\varepsilon_{z,t+1}$ and $\varepsilon_{x,t+1}$ are standard Gaussian shocks with contemporaneous correlation ρ_{xz}. In the specification above, x refers to the long-run risk component in productivity growth.

4.1.2 Firms

A firm in layer $j \in \{0, 1, ..., N\}$ hires labor $n_{j,t}$ from the household and owns capital stock $k_{j,t}$, which is layer-specific. The firms produce their output using constant returns to scale Cobb-Douglas production function over capital and labor, subject to aggregate productivity shock Z_t:

$$Y_{j,t} = Z_t k_{j,t}^{\alpha} n_{j,t}^{1-\alpha}, \quad j \in \{0, 1, .., N\}, \tag{4}$$

where α is the capital share of output for all firms. Since there are no capital suppliers for the top layer (layer N), its capital stock is assumed to be fixed over time ($k_{N,t} = k_{N,0}$). The capital stock for firms in layer $j \in \{0, .., N-1\}$ depreciates at rate δ, and evolves according to:

$$k_{j,t+1} = (1 - \delta + i_{j,t}) k_{j,t}, \tag{5}$$

where $i_{j,t}$ denotes the investment-rate of firm j. Each firm in layer $0 \geq j \geq N-1$ that wishes to invest amount $i_{j,t} k_{j,t}$, must purchase $\Phi(i_{j,t}) k_{j,t}$ units of its layer-specific capital goods directly from the layer vertically above it. Purchasing these layer-j capital goods is done under the equilibrium output price of layer $j-1$, P_{j-1}. The convex adjustment cost

function $\Phi(i)$ is given by:

$$\Phi(i) = \frac{1}{\phi}(1+i)^{\phi} - \frac{1}{\phi}. \tag{6}$$

In all, the period dividend of firm $j \in \{0, .., N-1\}$, $d_{j,t}$, is given by:

$$d_{j,t} = P_{j,t}Y_{j,t} - W_t n_{j,t} - P_{j+1,t}\Phi(i_{j,t})k_{j,t}, \tag{7}$$

where W_t denotes the real wage per unit of labor. Given that the top layer firm's capital is fixed, the dividend of the top layer firm is similarly given by $d_{N,t} = P_{N,t}Y_{N,t} - W_t n_{N,t}$.

Each firm chooses optimal hiring and investment (except for the top firm) to maximize its market value, taking as given wages W_t, output prices $P_{j,t}$, $\quad j \in \{0, .., N\}$, and the stochastic discount factor of the household $M_{t,t+1}$. Specifically, the layer-j representative firm maximizes:

$$V_{j,t} = \max_{\{n_{j,s}\ k_{j,s}\}} \quad E_t \sum_{s=t+1}^{\infty} M_{t,s} d_{j,s}, \tag{8}$$

subject to (5) if $j \in \{0, ..., N-1\}$.

4.1.3 Household

The economy is populated by a representative household. The household derives utility from an Epstein and Zin (1989) and Weil (1989) utility over a stream of consumption C_t:

$$U_t = \left[(1-\beta)C_t^{\frac{1-\gamma}{\theta}} + \beta(E_t U_{t+1}^{1-\gamma})^{\frac{1}{\theta}}\right]^{\frac{\theta}{1-\gamma}}, \tag{9}$$

where β is the subjective discount factor, γ is the risk aversion coefficient, and ψ is the elasticity of the intertemporal substitution (IES). For ease of notation, the parameter θ is defined as $\theta \equiv \frac{1-\gamma}{1-\frac{1}{\psi}}$. Note that when $\theta = 1$, that is, $\gamma = 1/\psi$, the recursive preferences collapse to the standard case of expected power utility, in which case the agent is indifferent to the timing of the resolution of the uncertainty of the consumption path. When risk aversion exceeds the reciprocal of IES ($\gamma > 1/\psi$), the agent prefers an early resolution of the uncertainty of consumption path, otherwise, the agent has a preference for a late resolution of the uncertainty.

The household supplies labor to all firms inelastically. It derives income from labor, as well as from the dividends of all $N+1$ production firms. The household chooses the layer-specific labor supply and consumption to maximize its lifetime utility, subject to the

following budget constraint:

$$\max_{C_s,\{n_{j,s},\theta_{j,s}\}_{j\in\{1..N\}}} U_t, \quad s.t. \quad P_{0,t}C_t + \sum_{j=0}^{N} \omega_{j,t+1}V_{j,t} = W_t \sum_{j=0}^{N} n_{j,t} + \sum_{j=0}^{N} \omega_{j,t}(V_{j,t} + d_{j,t}), \quad (10)$$

where $\omega_{j,t}$ is the share of the household in the ownership of the layer j firm. It is straightforward to show that the SDF used to discount the dividends of firms in all layers is given by:

$$M_{t,t+1} = \beta \left(\frac{C_{t+1}}{C_t}\right)^{\frac{-1}{\psi}} \left(\frac{U_{t+1}}{E_t\left[U_{t+1}^{1-\gamma}\right]^{\frac{1}{1-\gamma}}}\right)^{\frac{1}{\psi}-\gamma}. \quad (11)$$

4.1.4 Equilibrium

In equilibrium, wage W_t, and output prices $\{P_{j,t}\}_{j\in\{0,..,N\}}$, are set to clear all markets:

- Labor market clearing:

$$\sum_{j=0}^{N} n_{j,t} = 1. \quad (12)$$

- Differentiated capital-goods market clearing:

$$\Phi(i_{j-1,t})k_{j-1,t} = Y_{j,t}, \quad \forall j \in \{1, ..., N\}. \quad (13)$$

- Consumption-good market clearing:

$$C_t = Y_{0,t}. \quad (14)$$

- Firm-ownership market clearing:

$$\omega_{j,t} = 1, \quad \forall j \in \{0, ..., N\}. \quad (15)$$

An equilibrium consists of prices, labor, and capital allocations such that (i) taking prices and wages as given, the household's allocation solves (10), and firms' allocations solve (8); (ii) all markets clear.

4.2 Calibration

Table 5 shows the parameter choice of the model in the benchmark case. The model is calibrated at an annual frequency. There are two main parameter groups.

Production parameters. We set N to 5, implying 6 production layers, similarly to the benchmark empirical results. We set $\alpha = 0.33$, so that the labor share of output across

different layers is 2/3. The annual depreciation rate is 10%. The capital adjustment cost parameter ϕ helps to match the auto-correlation of output growth to the data, and boost the volatility of the equity premium. The aggregate productivity log-growth μ_z is set such that the steady state growth rate of consumption is about 2%, similarly to the data. We set σ_z to 1.7%, to obtain an annual volatility of consumption growth slightly below 2%, consistently with a long-run sample equivalent. To keep the long run component of consumption small, we impose ϕ_x to be 0.085. This is a conservative value. Croce (2014) shows that in the sample of 1930-2008, the ratio of the long-run risk volatility to the short-run risk volatility is roughly 10%. We set the persistence of the long-run component ρ_x to 0.98. This value is set to match the annual auto-correlation of consumption growth to the data (about 0.5). For simplicity, we set ρ_{xz} to 1. This reduces the number of shocks in the model to only one.

Table 5: **Model Calibration**

Symbol	Value	Parameter
Panel A: Production		
N	5	Number of layers
α	0.33	Share of capital in output
ϕ	25	Investment adjustment cost
δ	0.1	Depreciation rate
Panel B: Technology Shock		
μ_z	0.013	Productivity growth rate
σ_z	0.017	Short-run productivity shock volatility
ϕ_x	0.085	Ratio of Long-to-Short-run productivity volatility
ρ_x	0.98	Persistent of long-run productivity
ρ_{xz}	1	Correlation between short and long run productivity shocks
Panel C: Preferences		
β	0.98	Subjective discount factor
γ	10	Relative risk aversion
ψ	2	Intertemporal elasticity of substitution

The table shows the parameter values used in the benchmark model calibration. The model is calibrated at the annual frequency.

Preference parameters. We set the relative risk aversion and the intertemporal elasticity of substitution (IES) to 10 and 2, respectively. These are consistent with the estimates of Bansal et al. (2012), and Colacito and Croce (2011). We set the subjective time discount

factor to 0.98, to target the level of the real risk free rate.

4.3 Model Results

4.3.1 Vertical Position Model and Aggregate Moment Implications

The calibrated model is solved using a third-order perturbation method. The first order conditions, and the required detrending are shown in the Appendix.

Table 6 compares aggregate moments of macroeconomic and return variables implied by the model with their empirical counterparts. The model-implied moments are computed from a simulated population path. Panel A reports summary statistics for consumption, output, and investment growth rates. The growth rate of all macro quantities is roughly 2% per annum, consistently with the data. The volatility of consumption growth is 1.75% in the model versus 1.33% in the data. While the model-implied consumption volatility is somewhat larger than the data, it is still conservatively low, and consistent with a long-run sample estimate of consumption growth volatility.[11] The model implied volatility of output, 2.11%, falls inside the empirical 95% confidence interval.

Investment's volatility is larger than the volatility of consumption or output, in-line with the data, yet smaller than the data point estimate. This low volatility does not stem from the capital adjustment costs, but rather from the value-weighted aggregation method. The aggregate investment volatility is primarily driven by the low investment volatility of the largest layer, layer 0, which equals 3.05% per annum. Unlike a one-sector economy, in which output is used for both consumption and investment, in our model, the output of layer 0 is used for consumption purposes only. Keeping consumption volatility low restricts the cash-flow variability of layer 0 and also its investment volatility. Importantly, computing an equally-weighted average of investment growth rate across the different layers yields annual volatility of 5.05%, much closer to the data. The volatility of investment growth for layers 3, 4, and 5 are 5.39%, 6.10% and 6.43%, respectively. These model-implied estimates are very close to the empirical counterpart(s). Since the top layers have less capital, consistently

[11] In the period of 1930-2012, the volatility of consumption growth is 2.11%

with the data, the value-weighted aggregation scheme attenuates investment volatility. The autocorrelation of consumption and output are 0.45 and 0.30 in the model, respectively. These are strikingly close the empirical estimates of 0.52 and 0.28, for consumption and output growth. The autocorrelation of investment growth falls inside the empirical 95% confidence interval.

Table 6: **Aggregate Moments: Model versus Empirical Equivalents**

Variable and Statistic	Model	Data	
	Panel A. Macroeconomic Variables		
Consumption growth:			
Mean (%)	1.94	1.97	[1.58, 2.35]
Standard deviation (%)	1.75	1.33	[1.11, 1.67]
Autocorrelation	0.45	0.52	[0.29, 0.75]
Output growth:			
Mean (%)	1.94	2.11	[1.60, 2.61]
Standard deviation (%)	2.13	1.74	[1.45, 2.18]
Autocorrelation	0.30	0.28	[-0.04, 0.60]
Investment growth:			
Mean (%)	1.94	1.74	[-0.22, 3.70]
Standard deviation (%)	3.26	6.83	[5.69, 8.53]
Autocorrelation	0.13	0.32	[0.13, 0.52]
	Panel B. Return Variables		
Excess Market portfolio Return:			
Mean (%)	4.13	4.89	[-0.20, 9.97]
Standard deviation (%)	5.10	17.70	[14.76, 22.11]
Autocorrelation	-0.01	-0.04	[-0.29,0.21]
Risk-free rate:			
Mean (%)	1.02	1.04	[0.51, 1.57]
Standard deviation (%)	0.90	1.84	[1.54, 2.30]

The table shows annual moments from simulated model data against their empirical counterparts. Panel A presents moments related to macroeconomic variables, and Panel B related to aggregate asset-prices. The model implied moments are obtained from a simulated population path of length 100,000. The empirical moments are based on annual data of a modern sample, 1964-2012 (we adopt the term "modern" from Campbell et al. (2012)). Consumption, output and investment growth rates are real and per-capita. The market portfolio is measured using CRSP value weighted returns. The real risk free rate corresponds to a three month T-bill rate net of expected inflation. Brackets represent empirical 95% confidence-intervals.

The model also generates reasonable aggregate asset pricing moments. The equity premium in the model is levered by a factor of 5/3, to account for financial leverage. The

model-implied equity premium equals 4.13% per annum, close to the empirical counterpart of 4.89%. One dimension in which the model deviates from the empirical evidence is the volatility of the market excess return. It is difficult to generate a high equity premium and a high volatility of stock returns in a general equilibrium production model (see related discussion in Gomes et al. (2003)). The non-trivial equity premium is generated through a fairly high risk aversion of 10, along with a persistent growth-productivity component similar to Bansal and Yaron (2004) and Croce (2014). The risk-free rate is about 1% per annum in the model and the data, with a very conservative annual volatility of 1%. The elasticity of intertemporal substitution parameter intensifies the volatility of stock returns, while keeping the volatility of the risk-free rate low.

Table 7: **Vertical Position and Expected Return: Model versus Data**

	Model	Data	
Panel A. Excess returns by vertical position			
layer 5	16.07	16.49	[11.89, 21.08]
layer 4	12.01	11.86	[7.21, 16.52]
layer 3	9.42	7.39	[3.13, 11.64]
layer 2	7.15	6.40	[2.57, 10.23]
layer 1	5.23	5.27	[1.81, 8.73]
layer 0	3.59	5.22	[1.86, 8.58]
Panel B. Spreads			
spread (5-0)	12.49	11.27	[6.94, 15.60]
spread (5-4)	4.06	4.62	[0.12, 9.13]

The table presents excess returns and spreads in the model against their data equivalents. Panel A shows mean excess returns of firms at different vertical positions (layers). Panel B shows return spreads between layers 5 and 0 and layers 5 and 4. Layer 0 refers to the firm(s) that produce final consumption goods, while layer 5 refers to the firm(s) that produce capital goods in the top vertical position. The model excess returns are obtained from a simulated model path of length 100,000 years. The empirical excess returns are based on a monthly sample from 2003-11:2013-02, aggregated to form continuously-compounded annual return observations over the past 12 months. Brackets represent 95% confidence-intervals.

4.3.2 Vertical Position Model and Cross-Sectional Return Implications

Our main empirical contribution is to establish a novel spread based on vertical position. The excess return spread between the top (layer 5) and bottom layers (layer 0) is about 105 bps per month, or 11.27% per annum (continuously compounded). The production

economy successfully replicates this sizable spread, as seen in Table 7. The table reports the model-implied average excess return of the different layers, against the empirical estimates. The model-implied return spread between layers 5 and 0 is 12.49% per annum. The spread is impressively large and falls inside the empirical confidence interval. The model-implied mean excess returns increase monotonically from layer 0 to layer 5 for all layers and fall inside the 95% confidence interval of the data. Moreover, the model generates excess returns for layers 1, 4, and 5 that are strikingly similar to the data estimates.

4.3.3 Inspecting the Mechanism: the Role of Creative Destruction

The TMB Spread. The vertical position spread is driven by a single aggregate productivity shock. Naturally, a positive productivity shock increases future consumption in the model. Under the benchmark calibration, with a preference for an early resolution of uncertainty, this effect on consumption reduces the marginal utility of the consumer. As a result, the productivity shock has a positive market price of risk. Similarly, with an IES greater than unity, asset prices rise in response to a positive productivity shock (see Croce (2014)). The productivity betas of layers 0 to 5 are all positive and monotonically increasing. In other words, the sensitivity of the top layer to productivity innovations is larger than that of the bottom layer. Consequently, firms in a higher vertical position are *riskier* because their valuations are more cyclical. In good (bad) times, they appreciate (depreciate) more in value compared to firms in a lower vertical position. This happens as firms in the bottom production layers are more subject to creative (Schumpeterian) destruction.

Intuitively, a positive productivity shock increases the future cash flows for all firms due to higher *demand*, which starts at the consumer side and flows upwards in the production chain. Put differently, a positive productivity shock represents a higher value of *growth options* for all firms. However, the productivity shock has a differential effect on the value of *assets-in-place* for firms in different vertical positions. A positive productivity shock has a negative effect on the value of the assets-in-place because of the increased capital *supply*. The supply effect amplifies as it propagates downward in the production chain and destroys

existing value for firms in the bottom layers.

To understand the logic above, consider the following example of Nvidia and Amazon. Nvidia supplies Graphics Processing Units (GPUs) to Amazon for its Amazon Web Services (AWS), which is a cloud services platform. AWS uses Nvidia's GPUs to accelerate artificial intelligence and high performance computing workloads. For simplicity, assume that AWS is Amazon's only business and that Nvidia and Amazon are the only two firms in the supply chain. An economy wide technological improvement should appreciate the value of Nvidia more than the value of Amazon. While the technological advancement increases the dividends for both firms, it has a creative destruction effect only on Amazon. The value of the existing stock of GPUs deployed in Amazon's hyperscale data center, which is Amazon's installed capital, drops. The technological improvement means that it is easier to replace Amazon's assets-in-place because they are now cheaper to produce. However, this creative destruction argument does not apply to Nvidia, as it has no capital suppliers that can erode its existing capital stock. In the more general case, firms in the bottom production layers suffer more from this creative destruction. Technological improvements cause them to appreciate less compared to firms in the top layers, making their valuation less cyclical.

The creative destruction argument can be demonstrated using standard Q-theory. Table 8 reports the productivity elasticity of firms in layers 0 to 4. Because of the constant returns to scale and perfect competition, for these layers, Tobin's Q is a sufficient statistic for the ex-dividend firm value (or firms' production beta). The table shows that productivity shock affects the Tobin's Q of the top layers more strongly than that of bottom layers. An optimality condition for all layers stipulates that $Q_j = P_{j+1} \cdot \Phi'(i_j) \quad \forall j \in \{0..4\}$. The condition implies that the changes in Tobin's Q can be attributed to two separate channels: a change in the price of new capital (P_{j+1}) and a change in the capital installation costs (Φ'). Table 8 shows that a positive productivity shock increases the relative price of new capital goods less strongly for the bottom layers. Again, this is a result of a more pronounced creative destruction (or equivalently, a greater supply). In addition, productivity shocks

induce firms in the top layers to invest more compared to bottom layer firms. This is because the expected marginal revenue of capital is higher for the former firms, as their output price appreciates more. Thus, firms in the top of the production chain face greater capital installation costs (Φ'). This further enhances their Tobin's Q.

CAPM α. In the model, all layers of production are affected by a common productivity shock. Nevertheless, the CAPM does not hold in the model because valuations are non-linear in the underlying shock. The non-linear effect is explained by the decreasing returns to scale of the top layer, by the convex adjustment costs, and by the non-linear dependence between wages and productivity. We find that CAPM α in the model is 4 percent per year.[12] This suggests that the non-linear effects in the model are quantitatively large.

Table 8: **Model Implied Productivity-Elasticities By Vertical Position**

Layer j	$dlog(Q_j)/d\varepsilon_z$	$dlog(P_{j+1})/d\varepsilon_z$	$dlog(\phi'(i_j))/d\varepsilon_z$	$d(i_j)/d\varepsilon_z \times 10$
4	0.058	0.016	0.042	0.128
3	0.052	0.014	0.039	0.126
2	0.045	0.012	0.034	0.122
1	0.036	0.009	0.028	0.107
0	0.025	0.005	0.021	0.081

The table presents slope coefficients (b) of the following projection, using a simulated model path: $dY_{j,t} = const + b \cdot \varepsilon_t + error$, where $Y_{j,t}$ is a model-implied variable of interest of vertical layer j, and ε_t is the aggregate productivity shock. The first column shows the approptiate j layer index number. The variable Y_t is either layer-j's log Tobin's Q ($log(Q_j)$), the price of its capital input ($log(P_{j+1})$), the log installation cost of new capital ($log(\phi(i_j))$), or investment rate (i_j). All results are based on a simulated path of length 100,000 periods.

4.3.4 Sensitivity Analysis

The previous section illustrates that the economic force behind the vertical position spread is creative destruction. In this section, we demonstrate that the spread is *qualitatively* robust to most calibration parameters, but point out the quantitative importance of the various parameter choices.

[12]The conditional CAPM with time-varying betas holds in the model.

Table 9: **Sensitivity Analysis to Model Parameters**

	(1) Benchmark Calibration	(2) $\psi = 1.1$	(3) $\psi = 0.9$	(4) $\phi_x = 0$; $\sigma_z = 0.018$	(5) $\phi = 1$
Mean Excess Returns (%)					
Layer 5	16.07	3.84	1.20	0.87	19.12
Layer 4	12.01	3.43	1.34	0.86	5.71
Layer 3	9.42	3.17	1.40	0.84	4.23
Layer 2	7.15	2.91	1.50	0.82	3.52
Layer 1	5.23	2.64	1.64	0.78	2.61
Layer 0	3.59	2.36	1.86	0.72	1.43
Spread (5-0)	12.49	1.49	-0.66	0.15	17.69
Mean Returns (%)					
Equity premium	4.13	2.40	1.82	0.75	1.94
Risk-free rate	1.02	3.57	5.12	4.40	0.67
Consumption Growth					
Mean (%)	1.94	1.79	1.98	1.94	1.99
Standard deviation (%)	1.74	1.94	1.98	1.74	1.80

The table presents summary model results using different calibrations. The left most column shows the variable of interest. Column (1) presents results from the model under the benchmark calibration. Columns (2) - (5) present results from the model calibrated using the same parameters as in the benchmark case, other than the parameter(s) specified right below the column number. In Column (2) the IES parameter ψ is set to 1.1. In Column (3) the IES parameter is set to 0.9. In Column (4) the long-run volatility is set to zero, and the short-run volatility is raised to 1.8%. In Column (5) the capital adjustment cost parameter ϕ is set to 1 (no adjustment costs). All model implied moments are based on a simulated population path of length 100,000 periods.

The Role of IES. Column (1) of Table 9 shows the summary model statistics for a model with a similar calibration to the benchmark case, except that the IES parameter, ψ, is reduced to 1.1 (still above unity). The drop in the elasticity of intertemporal substitution parameter raises the level of the risk-free rate and drops the level of the equity premium compared to the benchmark case. Importantly, the spread in the returns between layers 5 and 0 is still positive, but diminished in magnitude to 1.49%. Intuitively, a lower IES implies that firms invest less in response to long-run productivity shocks, which attenuates their production betas. Column (2) of Table 9 shows the results when the IES parameter is dropped below the unity threshold to 0.9. In that case, the spread between the layer 5 and layer 0 excess returns turns negative. When the IES is less than unity, the income effect dominates the substitution effect. In response to a productivity news shock, the household desires to increase consumption strongly. The productivity news shock acts as a positive

demand shock for the final goods, but less so for intermediary goods (that is, saving or capital goods). This positive demand shock diverts resources to firms at bottom layers. The capital installation costs become larger for these bottom layer firms, increasing their Tobin's Q and making them more sensitive to productivity shocks.

The Role of Long-Run Productivity. In Column (3) of Table 9 we report the results when we shut down the long-run productivity news shocks ($\phi_x = 0$). To keep consumption growth volatility constant at the benchmark case level, we simultaneously raise the short-run production volatility σ_z to 1.8%. In this case, consumption growth volatility remains 1.74% per annum. However, the equity premium is diminished to only 0.75% per annum. Similarly, the spread between layer 5 and layer 0 remains positive, but it is only 0.15%. As shown in Croce (2014), short-run productivity shocks do not induce firms to invest largely enough, which implies positive yet low productivity betas. Introducing a persistent component to the growth of aggregate TFP, as in the benchmark case, causes firms to react more strongly to technology news (so long as the substitution effect dominates), and amplified betas.

The Role of Capital Adjustment Costs. Column (5) of Table 9 presents the results when there are no capital adjustment costs ($\phi = 1$). Without adjustment costs, the spread between layers 5 and 0 is still positive and very sizable. The creative destruction channel, which is the primary force behind our benchmark result, is independent of the degree of adjustment costs.

Notice that without adjustment costs, the TMB spread is higher than the benchmark calibration (17.69% versus 12.49%). The spread between layers 5 and layer 4 also increases, but it is now too large compared to the data. In addition, the average market excess return drops to only 1.94%, and the excess return for the firm in the bottom layers is counterfactually low (about 1.4%). To understand the intuition behind these findings, note that the lack of adjustment costs has a strong effect on the excess returns of layers 0 to 4, but not on the excess return of layer 5. Intuitively, by excluding adjustment costs, productivity shocks are absorbed in how much the firms in layers 0 to 4 invest (that is, in

quantities), as opposed to being absorbed in their prices. This allows the firms in layers 0 to 4 to smooth their dividends more easily, making their valuations less volatile. As a result, the risk premium of these firms drops. The firm in layer 5, however, does not possess depreciable capital (it does not invest), and therefore is largely unaffected by the absence of adjustment costs. Consequently, the TMB spread must rise.[13]

5 Monopolistic Competition

5.1 The Augmented Model

We now augment the benchmark model to accommodate monopolistic competition. Our goal is to study how firms' market power, which affects the degree of creative destruction, affects the TMB spread.

Aggregate productivity has the same dynamics as those described in Section 4.1.1. The household side of the economy is identical to that described in Section 4.1.3. Unlike the perfect competition model, we now assume that each production layer $j \in \{0..N\}$ is populated by a mass of differentiated intermediate good producers, indexed by $m \in [0, 1]$. The output of an intermediate good producer in layer j at time t of variety m is denoted by $y_{j,t}(m)$. Its output price is $p_{j,t}(m)$.

Aggregators. In each layer j, an aggregator converts the layer's intermediate goods into a final composite layer good, $Y_{j,t}$, using a CES production function:

$$Y_{j,t} = \left[\int_0^1 y_{j,t}(m)^{\frac{\mu_j - 1}{\mu_j}} dm\right]^{\frac{\mu_j}{\mu_j - 1}}, \tag{16}$$

when $\mu_j \to \infty$, the intermediate good producers of the j-th layer face perfect competition. For any finite μ_j, the intermediate good producers are not perfect substitutes, and they possess some amount of monopolistic power.

[13] Importantly, the amplification of the TMB spread in the absence of adjustment costs depends only partially on the fact that firm 5 does not invest. As long as layer 5 firm's labor share of output is larger than that of the layers below it, the argument still qualitatively holds.

The j^{th} layer aggregator faces perfect competition in the product market. It solves:

$$\max_{\{y_{j,t}(m)\}} \quad P_{j,t}Y_{j,t} - \int_0^1 p_{j,t}(m)y_{j,t}(m)dm$$

s.t equation (16).

The above implies that the price index in layer j is given by $P_{j,t} = \left[\int_0^1 p_{j,t}(m)^{1-\mu_j} dm\right]^{\frac{1}{1-\mu_j}}$. The demand schedule for each intermediate good producer in layer j of variety m is given by $\left[\frac{p_{j,t}(m)}{P_{j,t}}\right]^{-\mu_j} Y_{j,t}$.

The aggregator of layer $j \in \{1..N\}$ supplies capital to intermediate good producers in layer $j-1$. The aggregator at layer 0 sells its goods to the household for final consumption.

Intermediate good producers. The intermediate good producer in each layer j of variety m faces the same production technology and conditions as described in Section 4.1.2. It owns its capital stock, $k_{j,t}(m)$, which depreciates at rate δ, and it hires labor from the household. Now, however, it has an additional degree of freedom: the ability to optimally select its output price. Specifically, the period dividend of an intermediate good producer of variety m in layer $j \in \{0, .., N-1\}$, $d_{j,t}$, is given by:

$$d_{j,t}(m) = p_{j,t}(m)\left[\frac{p_{j,t}(m)}{P_{j,t}}\right]^{-\mu_j} Y_{j,t} - P_{j+1,t}\phi(i_{j,t}(m))k_{j,t}(m) - W_t \cdot n_{j,t}(m), \qquad (17)$$

where, as before, W_t denotes the real wage per unit of labor. Given that the top layer intermediate good producer's capital is fixed as before ($k_{N,t}(m) = k_{N,0}(m)$), their dividend is similarly given by $d_{N,t} = p_{N,t}(m)\left[\frac{p_{N,t}(m)}{P_{N,t}}\right]^{-\mu_N} Y_{N,t} - W_t n_{N,t}(m)$. Each intermediate good producer chooses its output price, optimal hiring, and investment (except producers at the top firm), to maximize its market value, taking as given wages W_t, its layer price index $P_{j,t}$, $\quad j \in \{0, .., N\}$, and the stochastic discount factor of the household $M_{t,t+1}$. Specifically,

layer-j firm maximizes:

$$V_{j,t}(m) = \max_{\{n_{j,s}(m) \text{ and } p_{j,s}(m) \text{ iff j}\in \{0,..,N\};k_{j,s}(m) \text{ iff j}\in \{0,..,N-1\}\}} E_t \sum_{s=t+1}^{\infty} M_{t,s} d_{j,s}(m); \quad (18)$$

s.t. (19)

$$\left[\frac{p_{j,t}(m)}{P_{j,t}}\right]^{-\mu_j} Y_{j,t} \leq Z_t k_{j,t}(m)^{\alpha} n_{j,t}(m)^{1-\alpha} \quad (20)$$

$$k_{j,t+1}(m) = (1 - \delta + i_{j,t}(m))k_{j,t}(m) \text{ if } j \in \{0, ..., N-1\}. \quad (21)$$

Market clearing and equilibrium. Compared to Section 4.1.4, the market clearing conditions of the labor markets, the capital goods markets, and the consumption good market are modified as follows:

$$\Sigma_{i=0}^{N} \int_0^1 n_{j,t}(m)dm = 1,$$

$$\int_0^1 \Phi(i_{j,t}(m))k_{j,t}(m)dm = Y_{j+1,t} \quad \forall j \in \{0, ..., N-1\},$$

$$C_t = Y_{0,t}.$$

All other market clearing conditions remain the same. Equilibrium consists of prices, labor, and capital allocations such that (i) taking prices and wages as given, the household's allocation solves (10), and firms' allocations solve (18); (ii) all markets clear; (iii) we are interested in a symmetric equilibrium in which $k_{j,t}(m) = k_{j,t}$, $n_{j,t}(m) = n_{j,t}$, and $p_{j,t}(m) = p_{j,t}$, for all $m \in [0, 1]$.

We calibrate the augmented model at an annual frequency using a calibration identical to that described in Section 4.2. We further impose that $\mu_j = \mu, \quad \forall j \in \{0, ..., N\}$. We are interested in varying the markup parameter μ, and considering the quantitative implications that it has on the spread. The main prediction of the augmented model is that the TMB spread is higher for firms that have high supply chain competition. To test this prediction, we need to develop a new measure of supply chain competition that accounts not only for the competition that each firm faces, but also the competition faced by its direct and indirect suppliers. We present this measure next.

5.2 A Measure of Supply Chain Competition

Besides the vertical position, we construct a competition measure at the supply chain level. These measures take into account not only a firm's own competition environment, but also the competition faced by its direct and indirect suppliers.

To derive this measure, we combine information about the production network structure with information about the number of competitors each firm has. The FactSet Revere relationships dataset allows us to identify each firm's competitors at any point in time. Firms can either report the list of their competitors directly or it can be inferred from their competitors' reports. We assume that competition relationships are undirected links, meaning that it is sufficient for only one firm to report the relationship. We observe 271,586 competition links in the database. We eliminate links that last less than 90 days and combine relationships where there is a gap of less than 90 days, with a gap defined as the number of days between the end of the previous relationship and the beginning of a new relationship between the same pair of firms.

Let $\mathbf{C}_t$ be an n by 1 column vector that measures the number of each firm's competitors in month t. While it could be a measure of competition, it does not account for competition at the supply chain level. We define a new measure of competition at the supply chain level as follows.

$$\hat{\mathbf{C}}_t = \mathbf{C}_t + \sum_{j=1}^{J} \bar{\mathbf{S}}_t^j \mathbf{C}_t, \tag{22}$$

where $\hat{\mathbf{C}}_t$ is an n by 1 column vector that measures each firm's supply chain competition, $\bar{\mathbf{S}}_t$ is a matrix of supplier-customer adjacency matrix normalized by the number of suppliers that each customer has. In other words, if $\mathbf{S}_t$ is an adjacency matrix of zeros and ones, such that $\mathbf{S}_t(i,j) = 1$ if firm j is a supplier to firm i and $\mathbf{S}_t(i,j) = 0$ otherwise, then $\bar{\mathbf{S}}_t(i,j) = \mathbf{S}_t(i,j) / \sum_{k=1}^{n_t} \mathbf{S}_t(i,k)$.

When $J = 1$, the supply chain competition measure for a firm is equal to the number of the firm's competitors plus the average number of its suppliers' competitors. For $J > 1$, not only is the average number of competitors of a firm's suppliers included in the formula, but

also the average number of competitors of the suppliers of the suppliers, of the suppliers of the suppliers of the suppliers, etc. In our benchmark specification we use $J = 5$, meaning that the competitiveness of indirect suppliers at distance five from the firm is accounted for in the supply chain competition measure of the firm.

The new measure allows us to split firms into two subsamples. In the high competition subsample we include firms that have an above median measure of supply chain competition within each layer of production. The second subsample includes firms with a below median measure of supply chain competition. The comparison between these two subsamples allows us to test augmented model's main predictions. Specifically, we compare both the TMB spread and the exposure to TFP shocks.

5.2.1 TMB spread: high competition vs. low competition

We first examine the TMB spreads in the high and low competition samples. To compare the empirical results to the model, we consider two choices for the value of μ: (i) a high competition calibration, $\mu = 100$, implying a markup of 1%, and (ii) a low competition calibration: $\mu = 3$, implying a markup of 33%. These numbers are consistent with the empirical estimates of markups (see, e.g., Bilbiie et al. (2012)). The results are reported in Table 10.

The TMB spread drops when firms have more market power, in both the model and the data. In the data, the spread is larger for the high competition subsample. The empirical spread for the high competition subsample is 10.19% per annum, while for the low competition subsample it is only 4.41%.[14] The model-implied spread for these two subsamples is qualitatively and quantitatively similar to the data. Specifically, for the high competition calibration, the model-implied TMB spread is 12.15%, while for the low competition calibration, the spread is 7.90%. In addition, in both the model and the data, the spread between the excess return of layers 5 and 4 is positive and sizable under the high

[14]In untabulated results, we confirm this finding using equally weighted portfolios. In fact, the difference between the high and low competition subsample spreads is more pronounced using equally weighted returns. For the high competition group, the spread is 15.6% per annum, while for the low competition group, it is merely 5.78%.

competition subsample, but much smaller, and even negative, under the low competition subsample. For all layers, and both subsamples the model-implied excess returns fall inside the empirical 95% confidence interval.

In both the high and low competition model calibrations, the mean excess return of the top layer is similar (about 16%). More generally, the results of Table 10 imply that higher monopolistic competition has a negligible effect on the productivity beta of the top layer, but the beta of the bottom layer becomes significantly more positive. There are two driving forces behind this result.

First, as discussed in Section 4.3.3, under perfect competition firms' valuations are pinned down by the cost of replacing their capital stock (Tobin's Q). Under monopolistic competition, however, valuations also depend on monopolistic rents. The benefits arising from technological improvements are not eroded completely by competition. For downstream firms that possess monopolistic power, a technological improvement decreases the cost of investment and increases their rents. This positive and enhanced cash-flow effect (equivalently, a rise in growth option valuations) operates against the negative Schumpeterian effect on the value of assets-in-place.

Second, higher markups for a firm's suppliers diminish the creative destruction effect that technological improvements have on its installed capital. When suppliers of a certain firm have a higher degree of monopolistic power, it has a rationing effect on the production of these suppliers. In other words, in response to a positive productivity shock, the suppliers of the firm increase their output less than under the perfect competition case. Consequently, the price of the firm's intermediate capital goods does not drop as much, and the marginal value of the firm's installed capital falls by a smaller degree.

We conclude that as a result of these two forces, the TMB spread declines when the firm, or its suppliers, have monopolistic power.

Table 10: **Vertical Position, Competition, and Expected Return: Augmented Model versus Data**

	High Competition			Low Competition		
	Model	Data		Model	Data	
Panel A. Excess returns by vertical position						
layer 5	16.14	15.25	[10.27, 20.23]	16.26	11.12	[3.65, 18.59]
layer 4	12.85	10.12	[5.62, 14.62]	16.13	12.71	[7.19, 18.23]
layer 3	10.31	9.28	[5.40, 13.17]	14.86	5.69	[0.96, 10.43]
layer 2	7.96	5.38	[2.06, 8.70]	13.03	8.32	[3.22, 13.42]
layer 1	5.86	4.69	[1.32, 8.06]	10.80	5.99	[1.40, 10.58]
layer 0	3.98	5.06	[1.89, 8.23]	8.37	6.71	[2.49, 10.92]
Panel B. Spreads						
spread (5-0)	12.15	10.19	[5.70, 14.68]	7.90	4.41	[-2.23, 11.05]
spread (5-4)	3.29	5.13	[1.33, 8.92]	0.13	-1.60	[-6.00, 2.80]

The table presents excess returns and spreads in the model against their data equivalents, for both high and low competition subsamples. Panel A shows mean excess returns of firms at different vertical positions (layers). Panel B shows return spreads across layers 5 and 0 and layers 5 and 4. Layer 0 refers to the firm(s) that produce final consumption goods, while layer 5 refers to the firm(s) that produce capital goods in the top vertical position. In the model the high competition results are based on a calibration in which $\mu = 100$ (implying a markup of 1%), while the low competition results are based on a calibration in which $\mu = 3$ (implying a markup of 33%). The model excess returns are obtained from a simulated model path of length 100,000 years. The empirical excess returns are based on a monthly sample from 2003-11:2013-02, aggregated to form annual observations over the past 12 months. Brackets represent 95% confidence-intervals. The empirical measure of competition is described in Section 5.2.

5.2.2 Exposure to TFP shocks: high competition vs low competition

One implication of our model is that firms in the upper layers are more exposed to TFP shocks than firms in the lower layers are, and that the difference in exposures between the top and the bottom layers is more pronounced in a more competitive economy. To test this prediction, we examine the sensitivity of each layer portfolio to the utilization-adjusted TFP shocks, available on a real-time basis on the website of the Federal Reserve Bank of San Francisco.[15] Since the data is at the quarterly frequency, we convert our monthly portfolio returns into quarterly returns and run the regressions at the quarterly level.

Table 11 reports results of the time series regressions of the quarterly portfolio returns

[15]The data is available at http://www.frbsf.org/economic-research/indicators-data/total-factor-productivity-tfp/. Compared to the crude estimates of the Solow residuals, the main advantage of this TFP series is that it makes an adjustment for variation in factor utilization, following the approach of Basu et al. (2006). See Fernald (2012) for more details.

Table 11: **Exposure of each layer to the TFP shocks**

Panel A. High competition sample

	TMB	layer 0	layer 1	layer 2	layer 3	layer 4	layer 5
ΔDTFP	2.180***	0.418	0.901*	0.744	0.684	1.288*	2.598***
	(3.22)	(0.88)	(1.78)	(1.39)	(1.13)	(1.78)	(3.12)
Constant	2.341	1.738	1.401	1.785	2.776*	3.304*	4.079*
	(1.36)	(1.44)	(1.09)	(1.31)	(1.80)	(1.80)	(1.92)
R^2	0.216	0.021	0.081	0.051	0.034	0.081	0.213

Panel B. Low competition sample

	TMB	layer 0	layer 1	layer 2	layer 3	layer 4	layer 5
ΔTFP	-1.382	0.667	0.758	1.070	0.776	0.683	-0.715
	(-1.29)	(1.18)	(1.10)	(1.39)	(1.00)	(0.79)	(-0.64)
Constant	2.882	2.185	2.402	2.931	2.388	4.310*	5.068*
	(1.05)	(1.52)	(1.37)	(1.50)	(1.21)	(1.95)	(1.78)
R^2	0.041	0.038	0.033	0.051	0.027	0.017	0.011

This table presents the exposure of each layer to the quarterly utilization-adjusted TFP shocks, ΔTFP, constructed following the method in Basu et al. (2006). Panel A reports results for the high competition sample, Panel B for the low competition sample. Significance at the 1, 5, and 10 percent levels are indicated by ***, **, and *, respectively.

on changes in utilization-adjusted TFP, ΔTFP. It shows a sharp difference between the high competition and low competition samples. In the high competition sample (Panel A), the loading on the ΔTFP increases substantially from the bottom layer (0.418) to the top (2.598). Consistent with the sharp increases in stock returns from layer 4 to layer 5, the increase in beta is also most pronounced between these two layers: a jump from 1.288 to 2.598. This provides strong support for our model. It shows that in a competitive environment, firms in the top layer are indeed more exposed to TFP shocks than are those in the bottom layer.

Most importantly, we document a positive TFP beta of 2.18 for the TMB portfolio in the high competition sample, significant at the 1% level. When we run the same regression of the TMB portfolio returns on the TFP shocks for the high competition sample in the model, we get a beta of 2.19, which is almost identical to the beta in the data.

In contrast, the pattern is much less clear in the low competition sample (Panel B). If anything, there is a hump shape in the betas across layers. This is consistent with our

augmented model, which shows that monopolist rents weaken the relation between the vertical position and the exposure to TFP shocks.

Taken together, the empirical tests we present in this section provide both qualitative and quantitative "out-of-sample" validation of our baseline model with perfect competition, and suggest that creative destruction is an important mechanism that drives the return spread between the top and bottom layers.

6 Robustness Checks

We conduct several robustness checks to confirm that the TMB spread is statistically and economically robust to alternative portfolio formation methods. The results are reported in Table 12. In columns (1) and (2) we sort firms into portfolio using their vertical positions at a lower frequency than in the benchmark case (once a quarter or a year, respectively). Lower frequency sorting results in an even higher spread, and improved t-statistics. This is not surprising given the persistence of the assigned vertical positions. In column (3) we sort firms into portfolios every month t, based on the vertical position computed at the end of month $t-4$, as opposed to $t-2$ in the benchmark implementation. This permits more time for the relationship information to be absorbed in stock priced (although the database is updated daily). Again, the results are materially unchanged.

In the benchmark case we define a firm's vertical position as the minimum number of links connecting it to (any) bottom layer firm. In column (4), we consider an alternative way to compute firm's vertical position. For robustness, we change the vertical position to be the median number of links connecting the firm to the bottom layer firms. The results show that the minimum vertical position measure is not crucial to obtain the spread. In column (4), the spread between layers 5 and 0 is 77 basis points, significant at 10 percent level, and the spread between layers 5 and 1 is 92 basis points significant at the 5% level.[16]

Lastly, in column (5) we reduce the number of layers to five. All firms with a vertical position of four or above are assigned to the top production layer. In this case, the top

[16]Top portfolio includes all firms with a median vertical position above eight.

minus bottom spread drops. This is consistent with the model prediction that also exhibits a smaller spread when the number of layers decreases. However, the spread between layer 5 and layer 1 is still statistically significant.

Table 12: **Empirical Robustness of the Top-Minus-Bottom Spread**

	(1) Quarterly Rebalance	(2) Annually Rebalance	(3) Four-month Lag	(4) Median Distance	(5) Five Layers
Panel A. Excess return by vertical position					
Layer 5	1.900	1.846	1.619	1.500	-
Layer 4	1.306	1.054	1.244	0.995	1.431
Layer 3	0.996	0.948	0.983	1.039	0.989
Layer 2	0.866	0.877	0.813	0.835	0.871
Layer 1	0.730	0.717	0.683	0.580	0.729
Layer 0	0.731	0.731	0.706	0.731	0.731
Panel B. Spreads					
Top-minus-Bottom	1.168**	1.115***	0.913*	0.769*	0.699*
	(2.48)	(2.94)	(1.97)	(1.77)	(1.73)
Top-minus-Layer 1	1.170***	1.129***	0.936**	0.920**	0.702**
	(2.63)	(3.39)	(2.18)	(2.26)	(2.04)

The table presents excess returns of the different layers (Panel A) and spreads (Panel B) under multiple alternative portfolio formation methods. The portfolios are constructed in an identical manner to the benchmark case, except the following changes: (1) sorting firms into portfolios based on their vertical position only at the end of each quarter only (at the end of March, June, September and December); (2) sorting firms into portfolios based on their vertical position only once a year (at the end of June); (3) sorting firms into portfolio every month t based on the vertical position computed in month $t-4$; (4) using the operator median instead of minimum in equation (1) for the vertical position assignment; and (5) grouping firms into only five layers, where the top layer includes all firms with vertical position of 4 or above. In columns (1)-(4), the top layer is layer five, and in column (5) it is layer 4. All numbers are in percentage units. T-statistics are reported in parentheses. Significance at the 1, 5, and 10 percent levels are indicated by ***, **, and *, respectively. The returns are computed from November 2003 to February 2013.

We further rule out alternative explanations for the spread. In untabulated results, we verify that the TMB spread is robust to additional factors and known return spreads. We find that the TMB spread is not driven by the intermediary capital risk factor (He et al. (2016) or the durable vs. non-durable spread (Gomes et al. (2009)).[17]

[17]We consider the intermediary capital risk factor in our robustness checks because according to Bigio and La'O (2016), tightening of financial constraints in the economy affects top layers more than the bottom layers. The intermediation capital factor captures this risk. We consider the durability spread to verify that the TMB spread is not driven by differences in durability across the layers of production.

Firms at different vertical positions could differ in their centrality. Ahern (2013) finds that industries with higher network centrality have higher returns. However, we find that firms in upper layers have significantly lower eigenvector centrality than those in lower layers, and that layer-1 firms have the highest centrality. Therefore, the TMB spread cannot be explained by centrality.

Firms that are further away from consumers in the production chain may be less familiar to investors and suffer from more severe information asymmetry. However, this hypothesis is not supported by empirical evidence. First, we find no significant difference in the bid-ask spread and the dispersion of earnings forecasts by financial analysts across layers. Second, we find that institutional ownership is actually lower in upper layers than in lower layers, suggesting that retail investors do not shy away from upstream firms.

We also verify that the spread exists when we use only producers of durable goods as the bottom layer. Moreover, if we use the materials sector as layer zero and construct the vertical position of other firms relative to this layer using customer-supplier relationships, the spread between the top and the bottom layers becomes negative, as expected.

7 Conclusion

We use the FactSet Revere database of supplier-customer relationships to measure the vertical position of each firm in the production network from 2003 to 2012. We sort firms into portfolios based on their vertical positions. A robust fact about the cross-section of stock-returns emerges when we compare the returns of these portfolios. Firms at a higher vertical position (further up from the producers of final goods) have higher expected returns. The spread between the top and the bottom layer portfolios is 105 basis points per month.

We provide a risk-based explanation of this new finding using a general equilibrium model with multiple layers of production. The calibrated model is able to generate the same monotonic pattern of returns as a function of the vertical position, and a quantitatively similar spread between the top and the bottom portfolios.

The model suggests that firms at a lower vertical position have less exposure to aggregate productivity shocks than firms at a higher vertical position. While all firms derive a direct benefit from improved productivity at all layers of production, the positive effect is attenuated by a devaluation of assets-in-place. Devaluation occurs as technological progress makes the replacement cost of existing capital units cheaper. Firms at the bottom of the production network experience the attenuation effect to the highest degree because their capital is composed of the capital goods produced by all layers above them in the production process. As each stage of the production process improves, the depreciation of the assets-in-place of the bottom layer firms reflects the collection of these improvements. In other words, the assets-in-place of the firms in lower vertical positions are more exposed to the force of creative destruction.

A strong empirical support for the Schumpeterian nature of the spread comes when we split the sample into firms with high and low supply chain competition. The spread is smaller for the sample of firms that belong to supply chains with higher markups. Under monopolistic power, the price of firms' capital input decreases less following a technological advancement, since monopolistic suppliers do not increase supply as much as competitive suppliers do. An augmented model with monopolistic competition provides a quantitative confirmation for the above intuition.

References

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., Tahbaz-Salehi, A., 2012. The network origins of aggregate fluctuations. Econometrica 80, 1977–2016.

Acemoglu, D., Ozdaglar, A., Tahbaz-Salehi, A., 2017. Microeconomic origins of macroeconomic tail risks. The American Economic Review 107, 54–108.

Aghion, P., Howitt, P., 1992. A model of growth through creative destruction. Econometrica 60, 323–351.

Ahern, K. R., 2013. Network centrality and the cross section of stock returns, working paper, University of South California.

Aobdia, D., Caskey, J., Ozel, N. B., 2014. Inter-industry network structure and the cross-predictability of earnings and stock returns. Review of Accounting Studies 19, 1191–1224.

Atalay, E., 2017. How important are sectoral shocks? American Economic Journal: Macroeconomics forthcoming.

Bansal, R., Kiku, D., Yaron, A., 2012. An empirical evaluation of the long-run risks model for asset prices. Critical Finance Review 1, 183–221.

Bansal, R., Yaron, A., 2004. Risks for the long run: A potential resolution of asset pricing puzzles. The Journal of Finance 59, 1481–1509.

Barrot, J.-N., Loualiche, E., Sauvagnat, J., 2016. The globalization risk premium, working paper, MIT.

Basu, S., Fernald, J. G., Kimball, M. S., 2006. Are technology improvements contractionary? American Economic Review 96, 1418–1448.

Belo, F., Lin, X., Bazdresch, S., 2014. Labor hiring, investment, and stock return predictability in the cross section. Journal of Political Economy 122, 129–177.

Berk, J. B., Green, R. C., Naik, V., 1999. Optimal investment, growth options, and security returns. The Journal of Finance 54, 1553–1607.

Bigio, S., La'O, J., 2016. Financial frictions in production networks, working paper, UCLA.

Bilbiie, F. O., Ghironi, F., Melitz, M. J., 2012. Endogenous entry, product variety, and business cycles. Journal of Political Economy 120, 304–345.

Boldrin, M., Christiano, L. J., Fisher, J. D., 2001. Habit persistence, asset returns, and the business cycle. American Economic Review 91, 149–166.

Branger, N., Konermann, P., Meinerding, C., Schlag, C., 2017. Equilibrium asset pricing in directed networks, working paper, University of Muenster.

Buraschi, A., Porchia, P., 2012. Dynamic networks and asset pricing, working paper, Imperial College London.

Caballero, R. J., 2008. creative destruction. In: Durlauf, S. N., Blume, L. E. (eds.), *The New Palgrave Dictionary of Economics*, Palgrave Macmillan, Basingstoke.

Caballero, R. J., Hammour, M. L., 1996. On the timing and efficiency of creative destruction. The Quarterly Journal of Economics 111, 805–852.

Campbell, J., Giglio, S., Polk, C., Turley, R., 2012. An intertemporal CAPM with stochastic volatility, NBER Working Paper No. 18411.

Carhart, M. M., 1997. On persistence in mutual fund performance. The Journal of finance 52, 57–82.

Carlson, M., Fisher, A., Giammarino, R., 2004. Corporate investment and asset price dynamics: implications for the cross-section of returns. The Journal of Finance 59, 2577–2603.

Chun, H., Kim, J.-W., Morck, R., Yeung, B., 2008. Creative destruction and firm-specific performance heterogeneity. Journal of Financial Economics 89, 109–135.

Cohen, L., Frazzini, A., 2008. Economic links and predictable returns. The Journal of Finance 63, 1977–2011.

Colacito, R., Croce, M. M., 2011. Risks for the long run and the real exchange rate. Journal of Political economy 119, 153–181.

Croce, M. M., 2014. Long-run productivity risk: A new hope for production-based asset pricing? Journal of Monetary Economics 66, 13–31.

Epstein, L. G., Zin, S. E., 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. Econometrica 57, 937–969.

Fama, E. F., French, K. R., 1993. Common risk factors in the returns on stocks and bonds. Journal of financial economics 33, 3–56.

Fama, E. F., French, K. R., 2015. A five-factor asset pricing model. Journal of Financial Economics 116, 1 – 22.

Fernald, J., 2012. A quarterly, utilization-adjusted series on total factor productivity, working paper, Federal Reserve Bank of San Francisco.

Garlappi, L., Song, Z., 2016. Capital utilization, market power, and the pricing of investment shocks. Journal of Financial Economics forthcoming.

Gârleanu, N., Kogan, L., Panageas, S., 2012. Displacement risk and asset returns. Journal of Financial Economics 105, 491–510.

Gofman, M., 2013. Profitability, trade credit and institutional structure of production, working Paper, University of Wisconsin.

Gomes, J. F., Kogan, L., Yogo, M., 2009. Durability of output and expected stock returns. Journal of Political Economy 117, 941–986.

Gomes, J. F., Kogan, L., Zhang, L., 2003. Equilibrium cross-section of returns. Journal of Political Economy 111, 693–732.

Hayek, F., 1935. Prices and Production (2nd edition). Augustus M. Kelly Publishers, New York.

He, Z., Kelly, B., Manela, A., 2016. Intermediary asset pricing: New evidence from many asset classes. Journal of Financial Economics forthcoming.

Herskovic, B., 2017. Networks in production: Asset pricing implications. Journal of Finance forthcoming.

Hobijn, B., Jovanovic, B., 2001. The information-technology revolution and the stock market: Evidence. American Economic Review 91, 1203–1220.

Hou, K., Xue, C., Zhang, L., 2015. Digesting anomalies: An investment approach. The Review of Financial Studies 28, 650–705.

Jermann, U. J., 2010. The equity premium implied by production. Journal of Financial Economics 98, 279–296.

Kennedy, D., Lloyd-Ellis, H., 2003. Animal spirits through creative destruction. The American Economic Review 93, 530–550.

Kogan, L., Papanikolaou, D., 2014. Growth opportunities, technology shocks, and asset prices. The Journal of Finance 69, 675–718.

Loualiche, E., 2016. Asset pricing with entry and imperfect competition, working paper, MIT.

Menzly, L., Ozbas, O., 2010. Market segmentation and cross-predictability of returns. The Journal of Finance 65, 1555–1580.

Ozdagli, A. K., Weber, M., 2017. Monetary policy through production networks: Evidence from the stock market, working paper, Federal Reserve Bank of Boston.

Papanikolaou, D., 2011. Investment shocks and asset prices. Journal of Political Economy 119, 639–685.

Rapach, D., Strauss, J., Tu, J., Zhou, G., 2015. Industry interdependencies and cross-industry return predictability, working Paper, Saint Louis University.

Richmond, R., 2015. Trade network centrality and the currency carry trade, working Paper, NYU.

Schumpeter, J. A., 1942. Capitalism, socialism, and democracy. Hamper & Brothers, New York.

Tallarini, T. D., 2000. Risk-sensitive real business cycles. Journal of Monetary Economics 45, 507–532.

Weil, P., 1989. The equity premium puzzle and the risk-free rate puzzle. Journal of Monetary Economics 24, 401–421.

Yang, F., 2013. Investment shocks and the commodity basis spread. Journal of Financial Economics 110, 164–184.

Zhang, L., 2005. The value premium. The Journal of Finance 60, 67–103.

8 Appendix

8.1 Equilibrium Conditions

For an economy with $N+1$ layers there are $5N+4$ endogenous variables, denoted by: $\{n_{j,t} \quad (\text{for } j \in \{0..N\}), k_{j,t}, q_j, i_j \quad (\text{for } j \in \{0..N-1\}), P_{j,t} \quad (\text{for } j \in \{1..N\}), W_t, C_t, M_t\}$. The first-order conditions are given by:

$$W_t = (1-\alpha) P_{j,t} Z_{j,t} k_{j,t}^{\alpha} n_{j,t}^{-\alpha} \quad \forall j \in \{0..N\}, \tag{23}$$

$$q_{j,t} = \Phi'(i_{j,t}) P_{j+1,t} \quad \forall j \in \{0..N-1\}, \tag{24}$$

$$q_{j,t} = E\left[M_{t,t+1}\left(P_{j,t+1} z_{j,t+1} \alpha k_{j,t+1}^{\alpha-1} n_{j,t+1}^{1-\alpha} - P_{j+1,t+1}\Phi(i_{j,t+1} + (1-\delta+i_{j,t+1}) q_{j,t+1}\right)\right] \quad \forall j \in \{0..N-1\}, \tag{25}$$

where the capital of the top layer N is fixed to unity. In total, there are there are $5N+4$ model equations: the above first-order conditions, along with $N+1$ labor market clearing equation (12), N capital markets clearing equations given by (13), consumption good clearing given by (14), N capital law of motions (5), and the household SDF (11). We normalize $P_{0,t} = 1$ as a numeraire.

8.2 Detrending

In this section we assume that each layer of production $j \in \{0..N\}$ is subject to a layer-specific productivity shock denoted by $Z_{j,t}$. We set N to 5, in-line with the benchmark calibration. We demonstrate how to detrend the model for this general case. In the private case in which all productivity shocks are perfectly correlated, as in the main text of this paper, the equations below still hold by replacing $Z_{j,t} = Z_t \quad \forall j \in \{0..N\}$.

Define capital trends as:

$$\tau_{k4,t} = Z_{5,t} \tag{26}$$

$$\tau_{k3,t} = Z_{4,t} Z_{5,t}^{\alpha} \tag{27}$$

$$\tau_{k2,t} = Z_{3,t} Z_{4,t}^{\alpha} Z_{5,t}^{\alpha^2} \tag{28}$$

$$\tau_{k1,t} = Z_{2,t} Z_{3,t}^{\alpha} Z_{4,t}^{\alpha^2} Z_{5,t}^{\alpha^3} \tag{29}$$

$$\tau_{k0,t} = Z_{1,t} Z_{2,t}^{\alpha} Z_{3,t}^{\alpha^2} Z_{4,t}^{\alpha^3} Z_{5,t}^{\alpha^4} \tag{30}$$

Let the price trends be:

$$\tau_{p5,t} = Z_{0,t} Z_{1,t}^{\alpha} Z_{2,t}^{\alpha^2} Z_{3,t}^{\alpha^3} Z_{4,t}^{\alpha^4} Z_{5,t}^{\alpha^5-1} \tag{31}$$

$$\tau_{p4,t} = Z_{0,t} Z_{1,t}^{\alpha} Z_{2,t}^{\alpha^2} Z_{3,t}^{\alpha^3} Z_{4,t}^{\alpha^4-1} Z_{5,t}^{\alpha^5-\alpha} \tag{32}$$

$$\tau_{p3,t} = Z_{0,t} Z_{1,t}^{\alpha} Z_{2,t}^{\alpha^2} Z_{3,t}^{\alpha^3-1} Z_{4,t}^{\alpha^4-\alpha} Z_{5,t}^{\alpha^5-\alpha^2} \tag{33}$$

$$\tau_{p2,t} = Z_{0,t} Z_{1,t}^{\alpha} Z_{2,t}^{\alpha^2-1} Z_{3,t}^{\alpha^3-\alpha} Z_{4,t}^{\alpha^4-\alpha^2} Z_{5,t}^{\alpha^5-\alpha^3} \tag{34}$$

$$\tau_{p1,t} = Z_{0,t} Z_{1,t}^{\alpha-1} Z_{2,t}^{\alpha^2-\alpha} Z_{3,t}^{\alpha^3-\alpha^2} Z_{4,t}^{\alpha^4-\alpha^3} Z_{5,t}^{\alpha^5-\alpha^4} \tag{35}$$

Lastly, the trend of final consumption goods is given by:

$$\tau_{c,t} = Z_{0,t} Z_{1,t}^{\alpha} Z_{2,t}^{\alpha^2} Z_{3,t}^{\alpha^3} Z_{4,t}^{\alpha^4} Z_{5,t}^{\alpha^5} \tag{36}$$

Covariance-stationary first-order conditions can be achieved by rescaling the non-stationary variables of the model as follows:

- Divide $k_{j,t}$ by $\tau_{kj,t-1}$, for $j \in \{0..4\}$.
- Divide $p_{j,t}$ and $q_{j-1,t}$ by $\tau_{pj,t-1}$, for $j \in \{1..5\}$.
- Divide c_t and W_t by $\tau_{c,t-1}$.

After plugging the rescaled variables in the first-order equations, the equilibrium conditions can be written using stationary quantities.

In the deterministic steady state, the growth rate in Tobin's Q is given by:

$$\Delta q_0 = \Delta Z_0 \Delta Z_1^{\alpha-1} \Delta Z_2^{\alpha^2-\alpha} \Delta Z_3^{\alpha^3-\alpha^2} \Delta Z_4^{\alpha^4-\alpha^3} \Delta Z_5^{\alpha^5-\alpha^4} \tag{37}$$

$$\Delta q_1 = \Delta Z_0 \Delta Z_1^{\alpha} \Delta Z_2^{\alpha^2-1} \Delta Z_3^{\alpha^3-\alpha} \Delta Z_4^{\alpha^4-\alpha^2} \Delta Z_5^{\alpha^5-\alpha^3} \tag{38}$$

$$\Delta q_2 = \Delta Z_0 \Delta Z_1^{\alpha} \Delta Z_2^{\alpha^2} \Delta Z_3^{\alpha^3-1} \Delta Z_4^{\alpha^4-\alpha} \Delta Z_5^{\alpha^5-\alpha^2} \tag{39}$$

$$\Delta q_3 = \Delta Z_0 \Delta Z_1^{\alpha} \Delta Z_2^{\alpha^2} Z_3^{\alpha^3} \Delta Z_4^{\alpha^4-1} \Delta Z_5^{\alpha^5-\alpha} \tag{40}$$

$$\Delta q_4 = \Delta Z_0 \Delta Z_1^{\alpha} \Delta Z_2^{\alpha^2} \Delta Z_3^{\alpha^3} \Delta Z_4^{\alpha^4} \Delta Z_5^{\alpha^5-1} \tag{41}$$

Equations (38) - (41) illustrate the effect of productivity shocks on the steady-state growth rate of the marginal value of assets in place for different production layers (without accounting for risk premia, of course). Notice that for these steady-state values, $\frac{\Delta q_k}{\partial \Delta Z_\ell} > 0$ whenever $k \leq \ell$.

Corollary. *A positive productivity shock from layer $k \in \{0..N\}$ increases (decreases) installed capital's value growth of layer $\ell \in \{0..N-1\}$ iff $k \leq \ell$ ($k > \ell$).*

The corollary above demonstrates the creative destruction argument at the steady-state. When a shock originates from a production layer below a given layer, the higher demand operates to appreciate the value of installed capital. However, when a shock originates from a production layer above a given layer (i.e. a supplier, or a supplier of suppliers), it triggers Schumpeterian destruction which erodes the value of existing capital stock. The creative destruction of a shock originating in layer k on the value growth of layer $\ell < k$ diminishes in the absolute distance between the layers $|k - \ell|$, at a constant rate of α, the capital share of output.

8.3 Online Appendix

Table 13: **A list of firms in the top layer**

Sector	Company Name
Energy	Apco Oil And Gas Intl Inc
	Boots & Coots Inc
	Carbo Ceramics Inc
	Eog Resources Inc
	Foundation Coal Holdings Inc
	International Coal Group Inc
	James River Coal Co
	Key Energy Services Inc
	Natural Gas Services Group
	Omega Navigation Ent Inc
Materials	American Vanguard Corp
	Gold Resource Corp
	Turquoise Hill Resources Ltd
	Vulcan Materials Co
Industrials	Agco Corp
	Alpha Pro Tech Ltd
	Cleantech Solutions Intl Inc
	Lincoln Electric Hldgs Inc
	Manitowoc Co
	Patriot Transn Holding Inc
Health Care	Altus Pharmaceuticals Inc
	I-Flow Corp
	Nmt Medical Inc
	Nucryst Pharmaceuticals Corp
	Orthovita Inc
	Pharmacyclics Inc
	Thoratec Corp
Information Technology	Blackbaud Inc
	Internet Patents Corp
	Magal Security Systems
Utilities	Exelon Corp
	National Fuel Gas Co
	Westar Energy Inc

This table shows the companies that are in the top layer for at least 18 months.