Winners and Losers: Creative Destruction and the Stock Market*

Leonid Kogan† Dimitris Papanikolaou‡
Noah Stoffman§

April 27, 2015

Abstract

We develop a general equilibrium model of asset prices in which the benefits of technological innovation are distributed asymmetrically. Financial market participants do not capture all the economic rents resulting from innovative activity, even when they own shares in innovating firms. Economic gains from innovation accrue partly to the innovators, who cannot sell claims on the rents their future ideas will generate. We show how the unequal distribution of gains from innovation can give rise to the salient empirical patterns in asset price behavior, including a high risk premium on the aggregate stock market, return comovement and average return differences among growth and value firms, and the failure of traditional representative-agent asset pricing models to account for these facts.

*An earlier version of this paper circulated under the title "Technological Innovation: Winners and Losers." We thank Hengjie Ai, Jaroslav Borovička, Kent Daniel, Carola Frydman, Lars Hansen, Erik Loualiche, Deborah Lucas, Stavros Panageas, Jonathan Parker, Monika Piazzesi, Amit Seru, Costis Skiadas, Annette Vissing-Jorgensen, Amir Yaron, and seminar participants at Berkeley, Chicago Booth, Boston University, CITE, Columbia GSB, Georgetown, University of Indiana, Jackson Hole Finance Conference, LBS, LSE, MIT Sloan, NBER AP, Northwestern University, Society of Economic Dynamics, Universidad de San Andres, USC, Washington University, University of Washington, Wharton, and WFA for valuable discussions and comments. Dimitris Papanikolaou thanks Amazon (AWS in Education Grant award), the Zell Center for Risk and the Jerome Kenney Fund for financial support. Leonid Kogan thanks J.P. Morgan for financial support.

†MIT Sloan School of Management and NBER, lkogan@mit.edu

‡Kellogg School of Management and NBER, d-papanikolaou@kellogg.northwestern.edu

§Kelley School of Business, nstoffma@indiana.edu

Technological innovation is arguably the main driver of economic growth in the long run. However, the economic value generated by new ideas is usually not shared equally. The popular press is rife with rags-to-riches stories of new entrepreneurs, whose net worth rose substantially as a result of their innovative ideas – for instance, during the tech boom period of the 1990's and 2000s'. In addition to the large wealth gains for successful innovators, technological progress can also create losses through creative destruction, as new technologies render old capital and processes obsolete. For instance, advances is communication technology have enabled ride-sharing companies, such as Uber, to displace traditional taxi companies.[1]

In this paper we show that the unequal sharing of gains and losses from technological innovation can give rise to the salient empirical patterns in asset price behavior, including a high risk premium on the aggregate stock market, return comovement and average return differences among growth and value firms, and the failure of traditional representative-agent asset pricing models to account for these facts. We build a tractable general equilibrium model in which the benefits of technological progress are distributed unevenly across investors and firms. Our model allows for two forms of technological progress. Some advances take the form of improvements in labor productivity, and are complementary to existing investments, while others are embodied in new vintages of capital.[2] The latter type of technological progress leads to more creative destruction, since old and new capital are substitutes.

A key feature of our model is that the market for new ideas is incomplete. Specifically, shareholders cannot appropriate all the economic rents generated by new technologies, even when those technologies are developed by the firms they are investing in. Our motivation for this market incompleteness is that ideas are a scarce resource, and generation of ideas relies heavily on human capital. As a result, innovators are able to capture a fraction of the economic rents that their ideas generate. The key friction is that potential innovators cannot sell claims to these future rents. This

[1] Uber was founded in 2009. As of December 2014, Uber is a privately-held company that is valued at $41 billion. Between December 2009 and February 2015, the value of Medallion Financial Corp. (NASDAQ: TAXI), a specialty finance company that originates, acquires, and services loans that finance taxicab medallions has dropped by more than 50% in value relative to the level of the NASDAQ index.

[2] Throughout the paper we refer to the first type of technological progress as disembodied, and the second type as embodied. Berndt (1990) gives the following two definitions: "Embodied technical progress refers to engineering design and performance advances that can only be embodied in new plant or equipment; older equipment cannot be made to function as economically as the new, unless a costly remodelling or retrofitting of equipment occurs," and "by contrast, disembodied technical progress refers to advances in knowledge that make more effective use of all inputs, including capital of each surviving vintage (not just the most recent vintage). In its pure form, disembodied technical progress proceeds independently of the vintage structure of the capital stock. The most common example of disembodied technical progress is perhaps the notion of learning curves, in which it has been found that for a wide variety of production processes and products, as cumulative experience and production increase, learning occurs which results in ever decreasing unit costs."

market incompleteness implies that technological progress has an asymmetric impact on household wealth. Most of the financial benefit from innovation accrues to a small fraction of the population, while the rest must bear the cost of creative destruction. By exposing households to idiosyncratic randomness in innovation outcomes, improvements in technology can thus reduce households indirect utility. This effect of technological progress is particularly strong when innovations are embodied in new capital goods. Further, this displacive effect on indirect utility is amplified because households also care about their consumption relative to the economy-wide average.

Displacement risk contributes to the equity risk premium and also leads to cross-sectional difference in asset returns. In our model, firms differ in their ability to acquire projects that implement new technologies. This difference in future growth opportunities implies that technological progress has a heterogeneous impact on the cross-section of asset returns. Firms with few existing projects, but many potential new investment opportunities, benefit from technological advances. By contrast, profits of firms that are heavily invested in old technologies and have few growth opportunities decline due to increased competitive pressure. If all firms had the same expected returns, firms with rich growth opportunities would be especially valuable as they would help offset potential utility losses brought on by technological improvements. This mechanism lowers the average returns on growth firms in equilibrium. Our model thus delivers cross-sectional differences in stock returns.[3]

We estimate the parameters of the model using indirect inference. The baseline model performs well at replicating the joint properties of aggregate consumption, investment, and asset returns. The model generates an aggregate consumption process with moderate low-frequency fluctuations that are consistent with the data, volatile equity returns, a high equity premium, and a low and stable risk-free rate. The model also replicates the observed differences in average returns between value and growth stocks and the failure of the Capital Asset Pricing Model (CAPM) to explain such cross-sectional differences. The above patterns have proven challenging to reproduce in a representative-agent general equilibrium model. These results depend on three features of our model – technology advances that are embodied in new capital goods, incomplete market for ideas, and preferences over relative consumption. Restricted version of the model that eliminate any one of

[3]Our model allows us to address two robust empirical patterns in the cross-section of asset returns – the so called value puzzle. First, firms with higher than average valuations–growth firms–experience lower than average future returns. These differences in average returns are economically large and comparable in magnitude to the equity premium. This finding has proven to be puzzling because growth firms are typically considered to be riskier and therefore should command higher average returns. Second, stock returns of firms with similar valuation ratios comove, even across industries. These common movements are typically unrelated to the firms' exposures to fluctuations in the overall market value. See Fama and French (1992, 1993) for more details.

these three features have difficulty replicating the empirical properties of asset returns – especially the cross-sectional differences between value and growth firms.

In addition to the features of the data that we target, we evaluate the model on its other testable predictions. Most of these predictions rely on the fluctuations in the share of economic value that is generated by new technologies – a quantity that is challenging to measure empirically. We construct an empirical estimate of this share using data on patents and stock returns following prior work by Kogan, Papanikolaou, Seru, and Stoffman (2012). Using patents has the unique advantage in that it allows us to infer the economic value of the underlying innovations over a narrow time window. Our identification assumption is that the stock market response to patent issues is proportional to the economic value of the new technologies they cover. Aggregating across all patents granted in a given year scaled by market capitalization provides an estimate for the share of economic value due to new technologies.[4]

We use this estimated share of value due to new technologies to test some of the predictions of the model. We find that rapid technological progress within an industry is associated with lower future profitability for low Tobin's Q firms relative to high-Q firms. Increases in the share of value of new technologies are associated with a decline in the ratio of the median to the mean of the consumption distribution across households. We verify that increases in the relative value of new technologies are associated with lower market returns, and higher returns for growth firms relative to value firms. We replicate these results in simulated data from the model; the empirical estimates are in most cases close to those implied by the model. We interpret these findings as providing support for the model's main mechanism.

Our work is related to asset pricing models with an explicit production process. Kogan and Papanikolaou (2012) summarize recent advances in this literature. These models broadly fall into three main categories. In the first category are many of the existing general equilibrium models, which feature a representative household and focus on replicating the empirical patterns of relatively smooth aggregate quantities, volatile stock market returns, and large risk premia.[5]

[4]The stock market reaction to patent grants is an underestimate of the value of the patent – since part of the information about the innovation may already be priced in. Further, our analysis misses the patents issued to private firms, as well as inventions that are not patented. Our analysis should thus be viewed as a joint test of the model and the assumption that movements in the economic value of patented innovations are representative of the fluctuations in the overall value of new inventions in the economy.

[5]An incomplete list includes Rouwenhorst (1995); Jermann (1998); Boldrin, Christiano, and Fisher (2001); Kaltenbrunner and Lochstoer (2010); Guvenen (2009); Campanale, Castro, and Clementi (2010); Papanikolaou (2011); Garleanu, Panageas, and Yu (2012); Croce (2014). Closest to our work is Papanikolaou (2011) and Garleanu et al. (2012), which also feature two types of technological progress: disembodied shocks that affect the productivity of all capital, and embodied (also termed 'investment-specific' shocks) that affect the productivity only of *new* capital. Papanikolaou (2011) focuses on the pricing of embodied shocks in an environment with complete markets, and argues

Since dividends and consumption are both endogenous, general equilibrium models face additional challenges in replicating the cross-sectional properties of asset returns. The second group of papers takes a different route and constructs partial equilibrium models with heterogeneous firms.[6] These papers analyze the economic determinants of cross-sectional differences in firms' systematic risk and risk premia – with a particular focus on understanding the origins of average return differences among value and growth firms. Our work embeds the structural model of the firm of Kogan and Papanikolaou (2013, 2014) in general equilibrium.

Our work fits into the third category, the papers developing general equilibrium models with non-trivial forms of heterogeneity. Gomes, Kogan, and Zhang (2003) study the cross-section of asset returns in a model where technology shocks are complementary to all capital. Pastor and Veronesi (2009) study the pricing of technology risk in a model with a life-cycle of endogenous technology adoption. Ai, Croce, and Li (2013) analyze the value premium in a model where some technology shocks only affect the productivity of old capital. All of the above are representative-agent models. Garleanu, Kogan, and Panageas (2012) model the value premium puzzle in a an overlapping-generations economy where technological improvements are embodied in new types of intermediate goods, and households face inter-generational displacement risk since existing agents cannot trade with the unborn. Similar to Garleanu et al. (2012), we also study the pricing of technology risk in incomplete markets, but in our model there is imperfect risk sharing among the *existing* population of households, which implies a different mechanism for how the technological shocks are priced and how one may test the model empirically. In addition, in our model there is a stationary distribution of firms, and firms invest and accumulate capital. This allows for a direct comparison of the model with the data. Finally, in this paper we systematically estimate model parameters using indirect inference.

The key part of our model's mechanism is that technological progress endogenously increases households' uninsurable consumption risk. The fact that time-varying cross-sectional dispersion of consumption can increase the volatility of the stochastic discount factor is well known (Constantinides and Duffie, 1996; Storesletten, Telmer, and Yaron, 2007; Constantinides and Ghosh, 2014). In our setting, time variation in households' uninsurable risk arises as an equilibrium effect of aggregate technological shocks. The resulting effect on asset prices is further amplified by households' preferences over relative consumption. Our work thus builds upon the long literature emphasizing

that they carry a negative risk premium due to short-run consumption fluctuations. Garleanu et al. (2012) focus on understanding the joint time-series properties of consumption and excess asset returns.

[6] A partial list of this class of models is Berk, Green, and Naik (1999); Carlson, Fisher, and Giammarino (2004); Zhang (2005); Kogan and Papanikolaou (2013, 2014).

the role of consumption externalities and relative wealth concerns for asset prices and equilibrium investment and consumption dynamics (Duesenberry, 1949; Abel, 1990; Gali, 1994; Roussanov, 2010). Closest to our work is Roussanov (2010), who argues that if households have preferences over their rank in the consumption distribution, in certain cases they are willing to invest in risky, zero-mean gambles whose payoff is uncorrelated with the aggregated state. In our setting, preferences over relative consumption induce agents to accept low risk premia (or equivalently high valuations) to hold assets that increase in value when technology prospects improve.

One of the best-documented empirical patterns in asset prices is the fact that average stock returns are positively correlated in the cross-section with the firms' ratio of market equity to book equity (M/B) (Fama and French, 1992). Quantitatively, the spread in returns between the high- and low-M/B firms is close to the equity premium, the difference in average returns between stocks and short-maturity bonds. An important component of the empirical puzzle is comovement among firms with similar M/B ratios (Fama and French, 1993, 1995; Hansen, Heaton, and Li, 2005). Specifically, stock returns, dividends and earnings of high-M/B firms tend to strongly comove with other high-M/B firms. The same is true for low-M/B firms. Further, this comovement is not driven by the heterogenous exposures of high- and low-M/B firms to the market portfolio, in other words, market betas do not fully summarize cross-sectional difference in systematic risk along the M/B dimension. Our model describes an economic mechanism that can give rise to the above empirical patterns.

1 The Model

We consider a dynamic continuous-time economy, with time indexed by t. We first introduce the productive sector of the economy – firms and projects they own, and the household sector. We next introduce households and the nature of market incompleteness in our model.

1.1 Firms and Technology

We emphasize one important dimension of heterogeneity among technological innovations. We model technological progress using two independent processes, x_t and ξ_t. The first process captures technological progress embodied in labor. Labor in our model is complementary to capital, and therefore growth in x raises productivity of *all* vintages of existing capital. In contrast, the process ξ affects the productivity embodied only in new vintages of capital, and different vintages of capital are substitutes.

There are two productive sectors in the model: a sector producing intermediate goods and a sector that aggregates these intermediate goods into the final good. The final good can be allocated either towards consumption or investment. In our analysis of asset prices, we focus on the intermediate-good firms. Firms in the final-good sector operate in a perfectly competitive environment with constant returns to scale, and therefore make zero profits in equilibrium.

1.1.1 Intermediate-good firms

Firms in the intermediate goods sector own and operate projects. Projects are introduced into the economy by the households we call *inventors*. Inventors initially own ideas, or the blueprints for creation of new projects. We assume that inventors lack the ability to implement their ideas on their own, and sell the blueprints for new projects to existing intermediate-good firms (we outline the details of blueprint sales below). Each firm f thus owns a constantly evolving portfolio of projects, which we denote by $\mathcal{J}_{ft}$. We assume that there is a continuum of infinitely lived firms in the economy, which we index by $f \in [0,1]$.[7] The set of all active projects in the economy is $\mathcal{J}_t = \bigcup_{f\in[0,1]} \mathcal{J}_{f,t}$.

Active projects

Projects are differentiated from each other by three characteristics: a) their operating scale, determined by the quantity of capital goods associated with the project, k; b) the systematic component of project productivity, ξ; and c) the idiosyncratic, or project-specific, component of productivity, u. Project j, created at time $\tau(j)$, produces a flow of output

$$\zeta_{j,t} = u_{j,t}\, e^{\xi_{\tau(j)}}\, k_{j,t}. \tag{1}$$

The systematic component of productivity ξ reflects technological progress *embodied* in new projects. It follows an arithmetic random walk

$$d\xi_t = \mu_\xi\, dt + \sigma_\xi\, dB_{\xi,t}, \tag{2}$$

where B_ξ is a standard Brownian motion. ξ_s denotes the level of frontier technology at time s. Growth in ξ affects only the output of new projects created using the latest frontier of technology.

[7] While we do not explicitly model entry and exit of firms, firms occasionally have zero projects, thus temporarily exiting the market. When such a firm finally acquires a project, it can be viewed as an entrant into the market. Investors can trade shares of firms with no active projects.

In this respect our model follows the standard vintage-capital model (Solow, 1960). An alternative modelling strategy capturing the same phenomenon would be to model ξ as investment-specific technological change – a separate technology shock that affects the productivity of the sector producing capital goods.

The level of project-specific productivity u_j is a stationary mean-reverting process that evolves according to

$$du_{j,t} \quad = \quad \kappa_u(1 - u_{j,t})\, dt + \sigma_u\, u_{j,t}\, dB^u_{j,t}, \tag{3}$$

where B^u_j are standard Brownian motions independent of B_ξ. We assume that $dB^u_{j,t} \cdot dB^u_{j',t} = dt$ if projects j and j' belong in the same firm f, and zero otherwise. As long as $2\,\kappa_u \geq \sigma_u^2$, the ergodic distribution of u has finite first two moments (see the Appendix for details). All new projects implemented at time s start at the long-run average level of idiosyncratic productivity, i.e., $u_{j,\tau(j)} = 1$. Thus, all projects created at a point in time are ex-ante identical in terms of productivity, but differ ex-post due to the project-specific shocks.

The firm chooses the initial operating scale k of a new project irreversibly at the time of its creation – firms cannot liquidate existing projects and recover their investment costs. Following project creation, the scale of the project diminishes according to

$$dk_{j,t} = -\delta\, k_{j,t}\, dt, \tag{4}$$

where δ is the economy-wide depreciation rate.

Firm investment opportunities – project arrival

Firms in the intermediate-goods sector are heterogeneous in their ability to acquire new projects. Firms acquire projects by randomly meeting inventors who supply blueprints. Project acquisition is exogenous to each firm, driven by a firm-specific doubly stochastic Poisson process $N_{f,t}$. The arrival rate of new projects equals $\lambda_{f,t}$. This arrival rate is time-varying and follows a two-state continuous-time Markov chain with high and low growth states $\{\lambda_H, \lambda_L\}$, $\lambda_H > \lambda_L$. The transition rate matrix S is given by

$$S = \begin{pmatrix} -\mu_L & \mu_L \\ \mu_H & -\mu_H \end{pmatrix}. \tag{5}$$

We denote the unconditional average of $\lambda_{f,t}$ by λ.

Implementing new projects

Creating a new project requires ideas and new capital goods. To implement a new idea as a project j at time t, a firm purchases new capital goods in quantity $I_{j,t}$. Investment in new projects is subject to decreasing returns to scale,

$$k_{j,t} = I_{j,t}^{\alpha}. \tag{6}$$

The parameter $\alpha \in (0, 1)$ effectively parameterizes convex adjustment costs in investment.

Aggregate output of intermediate goods

All intermediate goods are perfect substitutes; therefore we do not distinguish between the goods produced by different firms. The total output of the intermediate-good sector is an aggregate of output of all of the intermediate-good firms,

$$Z_t = \int_0^1 \left(\sum_{j \in \mathcal{J}_{f,t}} \zeta_{j,t} \right) df = \int_{J_t} e^{\xi_{s(j)}} \, u_{j,t} \, k_{j,t} \, dj. \tag{7}$$

The total output of the intermediate sector, Z_t, can be interpreted as a service flow from installed capital. At this stage, it is also helpful to define the stock of installed capital in the intermediate goods sector, adjusted for quality,

$$K_t = \int_{J_t} e^{\xi_{s(j)}} \, k_{j,t} \, dj. \tag{8}$$

1.1.2 Final-good firms

Firms in the final-good sector use intermediate goods and labor to produce the final good. The final good can be allocated towards consumption, or investment – the input that intermediate-good firms need to implement new projects. The production function for the final goods exhibits constant returns to scale,

$$Y_t = Z_t^{\phi} \left(e^{x_t} \, L_t \right)^{1-\phi}, \tag{9}$$

where Z_t and L_t are the output of the intermediate good sector and labor, respectively. Final good firms purchase intermediate goods Y and labor services L at the equilibrium prices p_Z and w, respectively. The final-good sector is not the main focus of our analysis. Under the production function (9) and perfect competition, final-good firms have zero market value in equilibrium.

The logarithm of the labor-augmenting productivity process x_t follows an arithmetic random

walk

$$dx_t = \mu_x \, dt + \sigma_x \, dB_{x,t}. \tag{10}$$

Here, B_x is a standard Brownian motion independent of all other productivity shocks. In particular, the productivity process x is independent from the embodied productivity process ξ.

The output of the final good sector can be allocated to either investment I_t or consumption C_t,

$$Y_t = I_t + C_t. \tag{11}$$

Our one-sector specification assumes that consumption can be transformed into investment one-to-one. Hence, the relative unit price of investment in our model is constant and equal to one.

Figure 1: Production

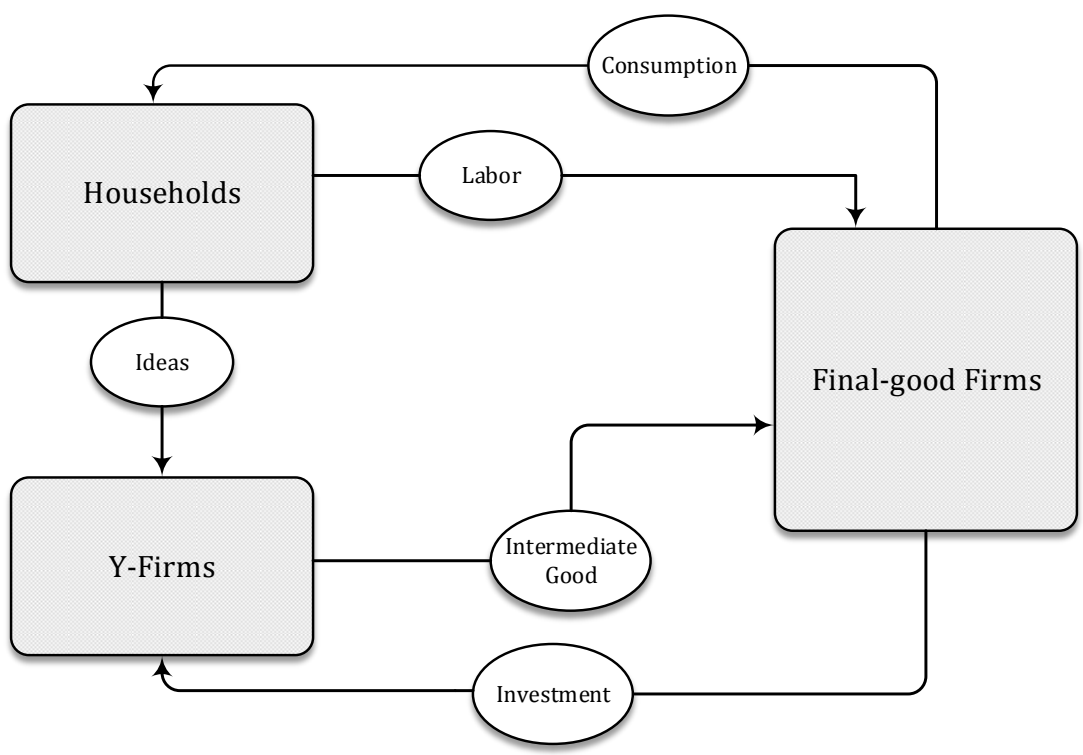

1.2 Households

There is a continuum of households, with the total measure of households normalized to one. Households die independently of each other, with the death being the first arrival of a Poisson

process with arrival rate δ^h. New households are born at the same rate, so the total measure of households remains constant. All households are endowed with the unit flow rate of labor services, which they supply inelastically to the firms producing the final good.

Households have access to financial markets, and optimize their life-time utility of consumption. Households are not subject to liquidity constraints; hence, they sell their future labor income streams and invest the proceeds in financial claims. We denote consumption of an individual household i by $C_{i,t}$.

All shareholder's have the same preferences, given by

$$J_t = \lim_{\tau \to \infty} \mathrm{E}_t \left[\int_t^\tau \phi(C_s, J_s; \bar{C}_s)\, ds \right], \tag{12}$$

where ϕ is the aggregator function:

$$\phi(C, J; \bar{C}) = \frac{\rho}{1-\theta^{-1}} \left(\frac{\left(C^{1-h} \left(C/\bar{C} \right)^h \right)^{1-\theta^{-1}}}{\left((1-\gamma) J \right)^{\frac{\gamma-\theta^{-1}}{1-\gamma}}} - (1-\gamma)\, J \right). \tag{13}$$

Households' preferences fall into the class of stochastic differential utility proposed by Duffie and Epstein (1992), which is a continuous-time analog of the preferences proposed by Epstein and Zin (1989). Relative to Duffie and Epstein (1992), our preference specification also incorporates a relative-consumption concern (otherwise termed as "keeping up with the Joneses", see, e.g., Abel, 1990). That is, households also derive utility from their consumption relative to aggregate consumption,

$$\bar{C} = \int_0^1 C_{n,t}\, dn. \tag{14}$$

The parameter h captures the strength of the relative consumption effect; γ is the coefficient of relative risk aversion; θ is the elasticity of intertemporal substitution (EIS); and ρ is the effective time-preference parameter, which includes the adjustment for the likelihood of death (see Garleanu and Panageas, 2014, for a model with random life spans and non-separable preferences).

The key feature of our model is imperfect risk sharing among investors. As we show below, embodied technology chocks affect the wealth shares of investors. The relative-consumption concern forces investors to care directly about fluctuations in their wealth shares. As a result, embodied technology shocks are particularly painful for the agents who benefit from them the least.

1.3 Household Innovation

Households are endowed with ideas, or blueprints, for new projects. Inventors do not implement these project ideas on their own and instead sell the ideas to firms producing intermediate goods. Inventors and firms bargain over the surplus created by new projects; the inventor captures a share η of the surplus.

Each household receives blueprints for new projects according to an idiosyncratic Poisson process with arrival rate μ_I. In the aggregate, households generate blueprints at the rate equal to the total measure of projects acquired by firms, λ. Not all innovating households receive the same measure of new blueprints. Each household i receives a measure of projects in proportion to her wealth $W_{i,t}$ – that is, equal to $\lambda W_{i,t} \left(\mu_I \int_0^1 W_{i,t}\, di\right)^{-1}$. Thus, conditional on innovating, wealthier households receive a larger measure of blueprints.

Importantly, households cannot trade in securities contingent on future successful individual innovation. That is, they sell claims against their proceeds from future innovations. This restriction on risk sharing plays a key role in our setting. In equilibrium, wealth creation from innovation leads to changes in the cross-sectional distribution of wealth and consumption, and therefore affects households' financial decisions.

1.4 Financial markets

We assume that agents can trade a complete set of state-contingent claims contingent on the paths of the aggregate and idiosyncratic productivity processes, as well as paths of project arrival rates and project arrival events at the firm level. We denote the equilibrium stochastic discount factor by π_t, so the time-t market value of a time-T cash flow X_T is given by

$$\mathrm{E}_t \left[\frac{\pi_T}{\pi_t} X_T \right]. \tag{15}$$

In addition, we follow Blanchard (1985) and assume that investors have access to competitive annuity markets that allow them to hedge their mortality risk. Specifically, this implies that conditional on surviving during the interval $[t,\, t + dt]$, investor j collects additional income proportional to her wealth, $\delta^h\, W_{j,t}\, dt$.

1.5 Discussion of the model's assumptions

Most extant production-economy general equilibrium models of asset returns build on the neoclassical growth framework. We depart from this literature in three significant ways.

a. Technological progress is embodied in new capital vintages. Most models assume that technological progress is complementary to the entire existing stock of capital, as is the case for the x shock in our model. However, many technological advances are embodied in new capital goods and benefit only the future vintages of capital investments. Several empirical studies show substantial vintage effects in plant productivity. For instance, Jensen, McGuckin, and Stiroh (2001) find that the 1992 cohort of new plants was 50% more productive than the 1967 cohort in its entry year, controlling for industry-wide factors and input differences.[8] To understand how technological change affects asset returns, we distinguish between the embodied and disembodied technological progress.

b. Incomplete markets for innovation. In our model, inventors generate new ideas and sell them to firms. We can map this stylized process into several types of innovation processes in the data. One possibility is that inventors work for existing firms, generate ideas, and receive compensation commensurate with the economic value of their ideas. Here, by inventors we mean highly skilled research personnel or corporate executives that can generate ideas for new investments. Since their talent is in scarce supply, we expect them to capture a significant fraction of the economic value of their ideas. Another possibility is that inventors implement the ideas themselves, creating startups that are partly funded by outside investors. Innovators can then sell their share of these startups to existing firms and thus capture a substantial share of the economic value of their innovations.

An important assumption in our model is that the economic value that is generated by new ideas cannot be fully pledged to outside investors. This assumption can be motivated on theoretical grounds. New ideas are the product of human capital, which is inalienable. Hart and Moore (1994) show that the inalienability of human capital limits the amount of external finance that can be raised by new ventures. Bolton, Wang, and Yang (2015) characterize a dynamic optimal contract between a risk averse entrepreneur with risky inalienable human capital and firm investors. The optimal contract involves a trade-off between risk sharing and incentives, and leaves the entrepreneur with a significant fraction of the upside gains.

c. Preferences over relative consumption. Our work deviates from many existing general equilibrium models by allowing households' utility to depend not only on their own consumption,

[8] The literature on the effects of embodied technological progress on growth and fluctuations includes, for instance, Solow (1960); Cooley, Greenwood, and Yorukoglu (1997); Greenwood, Hercowitz, and Krusell (1997); Fisher (2006).

but also on their consumption relative to the average per capita consumption in the economy. This assumption can be justified on empirical grounds. For instance, several studies document that, controlling for household income, the rank in the income distribution or the income of a peer group, is negatively related to self-reported measures of happiness (Clark and Oswald, 1996; Solnick and Hemenway, 1998; Ferrer-i Carbonell, 2005; Luttmer, 2005). These relative income concerns are substantial; the point estimates in Luttmer (2005), for instance, imply that income of households in the same metropolitan area is more important for happiness than the households' own level of income. Frydman (2015) finds strong evidence for utility preferences over relative wealth in an experimental setting using neural data collected through fMRI.

Recent theoretical work justifies preferences over relative consumption as a reduced-form description of individual behavior. Rayo and Becker (2007) propose a theory in which peer comparisons are an integral part of the "happiness function" as a result of an evolutionary process. DeMarzo, Kaniel, and Kremer (2008) show that competition over scarce resources can make agents' utilities dependent on the wealth of their cohort and induce relative wealth concerns.

In Section 3.3 we explore the sensitivity of our quantitative results to assumptions (a) to (c). In addition to these three main assumptions, our model deviates in some other respects from the neoclassical framework. These deviations make the model tractable but do not drive our main results.

First, we assume that projects arrive independently of the firms' own past decisions, and firms incur convex adjustment costs *at the project level*. Together, these assumptions ensure that the optimal investment decision can be formulated as a static problem, thus implying that the cross-sectional distribution of firm size does not affect equilibrium aggregate quantities and prices. In the aggregate, the cost of investment is then a convex function of investment level, as in Abel (1983). Second, the assumption that innovating households receive a measure of projects that is proportional to their existing wealth guarantees that the growth of household's wealth due to innovation does not depend on its wealth level. Together with homothetic preferences, this ensures that the households' optimal consumption and portfolio choices scale in proportion to their wealth, and thus the cross-sectional distribution of household wealth does not affect equilibrium prices.[9] Third, households in our model have finite lives. This assumption has no substantive effect on our main results and simply ensures that there exists a stationary distribution of wealth among

[9]This assumption likely weakens our main results compared to the case where all households have identical prospects of innovating. In the latter case, wealthier households would benefit less from innovation, raising their exposure to innovation shocks relative to our current specification.

households. Fourth, our assumption that project productivity shocks are perfectly correlated at the firm level ensures that the firm state vector is low-dimensional. Last, there is no cross-sectional heterogeneity among the quality of different projects. We could easily allow for an idiosyncratic component to ξ, perhaps allowing for substantial skewness to capture the notion that the distribution of profitability of new ideas can be highly asymmetric. Our conjecture is that such an extension would strengthen our main results by raising the level of idiosyncratic risk of individual households' consumption processes.

1.6 Competitive equilibrium

Here, we describe the competitive equilibrium of our model. Our equilibrium definition is standard, and is summarized below. After defining the equilibrium, we proceed by characterizing key equilibrium relations that are important for understanding the main mechanism and how it relates to the data.

Definition 1 (Competitive Equilibrium) *The competitive equilibrium is a sequence of quantities $\{C_t, I_t, Y_t, Z_t, K_t\}$; prices $\{\pi_t, p_{Z,t}, \mathrm{w}_t\}$; household consumption decisions $\{C_{i,t}\}$; and firm investment decisions $\{I_{j,t}\}$ such that given the sequence of stochastic shocks $\{x_t, \xi_t, u_{j,t}, N_{f,t}\}$, $j \in \mathcal{J}_t$, $f \in [0,1]$: i) households choose consumption and savings plans to maximize their utility (12); ii) household budget constraints are satisfied; iii) firms maximize profits; iv) the labor market clears, $L_t = 1$; v) the resource constraint (9) is satisfied; vi) the demand for new investment equals supply, $\int_0^1 I_{n,t} dn = I_t$; and vii) the market for consumption clears $\int_0^1 C_{n,t}\, dn = C_t$.*

Because of market incompleteness, standard aggregation results do not apply. Specifically, there are two dimensions of heterogeneity in the model: on the supply side, among firms; and on the demand side, among households. Both of these can potentially make the state space of the model infinite-dimensional. In our case however, we can solve for aggregate-level quantities and prices in Definition 1 as functions of the low-dimensional Markov aggregate state vector $X_t = (x_t,\, \xi_t,\, K_t)$. Specifically, the first moment of the cross-section distribution of installed capital K summarizes all the information about the cross-section of firms relevant for the aggregate dynamics in the model. The first two assumptions discussed above in Section 1.5 enable the relatively simple characterization of equilibrium.

Further, as we show in the Appendix, the following two variables are a sufficient statistic for the aggregate state,

$$\omega_t = \xi_t + \alpha\, \chi_t - \log K_t \tag{16}$$

and

$$\chi_t = \frac{1-\phi}{1-\alpha\,\phi}\,x_t + \frac{\phi}{1-\alpha\,\phi}\xi_t. \tag{17}$$

For instance, we can write aggregate output (9) as

$$Y_t = e^{\chi_t} e^{-\phi\,\omega_t}. \tag{18}$$

The variable χ_t is a random walk, and captures the permanent effect of the two technology shocks on aggregate quantities. The variable ω represents deviations of the current capital stock from its target level – and thus deviations of Y from its stochastic trend.

We solve for equilibrium prices and quantities numerically.[10] We next describe how we calibrate the parameters of the model.

2 Estimation

Here, we describe how we calibrate the model to the data. In Section 2.1 we describe which features of the data help identify the model's parameters. In Section 2.2 we discuss how we choose parameters through a minimum-distance criterion. In Section 2.3 examine the model's performance in matching the features of the data we target and the resulting parameter estimates.

2.1 Data

The model has a total of 20 parameters. Only a handful of these parameters can be calibrated using a priori evidence. We calibrate $\phi = 1/3$ so that the labor share in the production of the final good equals two thirds. We choose the probability of household death as $\delta^h = 1/40$ to guarantee an average working life of 40 years. We calibrate the probability of repeat innovation to equal $\mu_I = 0.15\%$ per year, in order to generate a consumption share of the top 1% of households to be

[10] Markets are incomplete in our model, therefore we cannot rely on the standard results for uniqueness and existence of equilibrium. In principle, there may exist multiple equilibria in our model. Our numerical solution characterizes one particular equilibrium. Along these lines, we cannot formally prove that ω_t is a stationary process in equilibrium. Instead, we verify that ω_t is stationary under our numerical solution.

equal to 14.3%.[11] Last, we create returns to equity by levering financial wealth by a factor of 2.5.[12]

We estimate the remaining 17 parameters using indirect inference. To do so, we first need to define a distance metric between the model and the data. We summarize the relevant information in the data in terms of 21 statistics, reported in the first column of Table 1. Due to data availability, each of these statistics is available for different parts of the sample. We use the longest available sample to compute these statistics. We relegate all details on the construction of these variables in Appendix B. Some of these statistics are standard, for example, the mean and volatility of aggregate quantities, and moments of the market portfolio and the risk-free rate. Dividends are not well defined in our model, hence we focus on net payout instead. Net payout can be potentially negative – depending on the parametrization; therefore we target the volatility of the ratio of net payout to book assets.

We add several statistics that are revealing of the mechanisms in our paper. First, technological progress leads to low frequency fluctuations in consumption and output in our model – see, for instance, equation (18). Hence, we also include as a target the estimate of the long-run standard deviation using the methodology of Dew-Becker (2014). Second, a large fraction of the parameters of the model govern the behavior of individual firms. We thus target the cross-sectional dispersion and persistence in firm investment, Tobin's Q, and profitability. Last, as we discuss in detail in Section 3, our model has implications for the cross-section of asset returns. We thus include as estimation targets the moments of a portfolio of value minus growth firms, where value and growth are defined according to firms' book-to-market ratio (following Fama and French, 1992).

2.2 Methodology

We estimate the parameter vector p using the simulated minimum distance method (Ingram and Lee, 1991). Denote by X the vector of target statistics in the data and by $\mathcal{X}(p)$ the corresponding

[11] We choose this parameter ex-post, that is, after the estimation of the other parameters, to match the estimated mean income share of the top 1% using the data of Piketty and Saez (2003). Given the other estimates, varying μ_I to 4.11% (the fraction of households that transition into entrepreneurship each year reported by Hurst and Lusardi (2004), a likely upper bound) has a quantitatively negligible effect on the moments of aggregate quantities and asset returns. The reason is risk aversion: given moderate amounts of utility curvature, the certainty equivalent of a bet that pays with probability μ_I an amount that is proportional to $1/\mu_I$ is negligible for small μ_I. Including μ_I in the full estimation, along with the inequality moment is infeasible given the computational cost of estimating inequality in the model.

[12] This value lies between the estimates of the financial leverage of the corporate sector in Rauh and Sufi (2011) (which is equal to 2) and the values used in Abel (1999) and Bansal and Yaron (2004) (2.74-3). We use a higher leverage parameter than Rauh and Sufi (2011) to account for the effects of operating leverage.

statistics generated by the model given parameters p, computed as

$$\mathcal{X}(p) = \frac{1}{S}\sum_{i=1}^{S} \hat{X}_i(p), \tag{19}$$

where $\hat{X}_i(p)$ is the 21×1 vector of statistics computed in one simulation of the model. We simulate the model at a weekly frequency, and time-aggregate the data to form annual observations. Each simulation has 1,000 firms. For each simulation i we first simulate 100 years of data as 'burn-in' to remove the dependence on initial values. We then use the remaining part of that sample, which is chosen to match length of our sample; each statistic is computed using the same part of the sample as its empirical counterpart. In each iteration we simulate $S = 100$ samples, and simulate pseudo-random variables using the same seed in each iteration.

Our estimate of the parameter vector is given by

$$\hat{p} = \arg\min_{p \in \mathcal{P}} (X - \mathcal{X}(p))' W (X - \mathcal{X}(p)). \tag{20}$$

To avoid scale effects, we specify the weighting matrix as $W = diag(XX')^{-1}$. Our choice of weighting matrix ensures that the estimation method penalizes proportional deviations of the model statistics to their empirical counterparts.[13] The model has a large number of parameters, and solving each iteration of the model is computationally costly. Thus, computing (20) using standard methods is infeasible. We therefore compute the minimum in (20) using the Radial Basis Function (RBF) algorithm in Björkman and Holmström (2000).[14]

[13] We choose this scaling rather than scaling by the inverse of the sample covariance matrix of X because the latter is infeasible to compute; many of the statistics in X are computed from different datasets (e.g. cross-sectional vs time-series, and often we do not have access to the underlying data. Also, not all of these statistics are moments, hence even if we had access to the original data, computing the covariance matrix of these estimates would be challenging. Hence our procedure provides a consistent but inefficient estimate of p. We have experimented with the following alternatives: i) use the diagonal of the sample covariance matrix and only include in X the statistics for which we can compute standard errors; ii) approximate the sample covariance matrix of X with the covariance matrix of $\hat{X}_i(p)$ across simulations as in (22). However, both of these cases force the model to match moments that are precisely estimated but economically less interesting, such as the dispersion in firm profitability. Further, when using $V_X^{-1}(\hat{p})$ as the weighting matrix, the algorithm is not always numerically stable.

[14] The RBF algorithm first fits a response surface to data by evaluating the objective function at a few points. Then, it searches for a minimum by balancing between local and global search in an iterative fashion. See Björkman and Holmström (2000) for more details. We use a commercial implementation of the RBF algorithm which is available through the TOMLAB optimization package.

We compute quasi-standard errors for the vector parameter estimates $\hat{p}$ as

$$V(\hat{p}) = \left(1 + \frac{1}{S}\right)\left(\frac{\partial}{\partial p}\mathcal{X}(p)'W\frac{\partial}{\partial p}\mathcal{X}(p)\right)^{-1}\frac{\partial}{\partial p}\mathcal{X}(p)'W'V_X(\hat{p})W\frac{\partial}{\partial p}\mathcal{X}(p)\left(\frac{\partial}{\partial p}\mathcal{X}(p)'W\frac{\partial}{\partial p}\mathcal{X}(p)\right)^{-1}, \tag{21}$$

where

$$V_X(\hat{p}) = \frac{1}{S}\sum_{i=1}^{S}(\hat{X}_i(\hat{p}) - \mathcal{X}(\hat{p}))(\hat{X}_i(\hat{p}) - \mathcal{X}(\hat{p}))', \tag{22}$$

is the estimate of the sampling variation of the statistics in X computed across simulations. We term the standard errors in (21) as quasi-standard errors because they are computed using the sampling variation of the target statistics across simulations, rather than their empirical counterparts. Under the null of the model, these two estimates coincide. If the model is misspecified, (22) need not be a good estimate of the true covariance matrix of X. Since our purpose in computing these standard errors is to assess the sensitivity of the criterion function to changes in the structural parameters, rather than obtaining a measure of precision on how these parameters are estimated, this is not a major concern.

2.3 Estimation results

Examining columns two to five of Table 1 we see that the baseline model performs reasonably well in fitting the data. The model generates realistic patterns for aggregate consumption, investment and corporate payout. In terms of the cross-section of firms, the model largely replicates the empirical cross-sectional dispersion in Tobin's Q and profitability, the low persistence in firm investment rates, and the weak empirical relation between firm investment and Q. At the same time, the model replicates key features of asset returns. The model generates a high equity premium, low and stable risk free rate, the value premium, the value factor, and the failure of the Capital Asset Pricing Model (CAPM). Consistent with the data, value firms in the model have higher average returns than growth firms, yet this difference in risk premia is not accounted by their differential exposure to the market portfolio.

In general, most of the statistics in simulated data are close to their empirical counterparts – however there are some exceptions. The model's equilibrium consumption process has somewhat more long-run consumption volatility than its empirical counterpart (0.056 vs 0.041), although the empirical value still falls within the model's 90% confidence interval. More importantly, the model generates stock returns that are smoother than their empirical counterparts by approximately a third. As is the case in almost all general equilibrium models, the need to match the relatively

smooth dynamics of aggregate quantities imposes tight constraints on σ_x and σ_ξ. Mechanisms that lead to time-variation in risk premia, for instance, time-variation in the volatility of the shocks σ_x and σ_ξ, may help increase the realized variation in asset returns. Due to the associated increased computational complexity, we leave such extensions to future research.

Last, the model has difficulty replicating the cross-sectional dispersion in Tobin's Q observed in the data; in the model, the inter-quartile range in Q is 34% smaller than the data. Given the fact that part of this dispersion may be measurement error – the replacement cost of capital is imperfectly measured in the data – this under-performance is not a major concern.

We report the estimated parameters in Table 2, along with their standard errors. The estimated utility curvature parameter is high, ($\hat{\gamma} = 106$). The parameter controlling the elasticity of intertemporal substitution ($\hat{\theta} = 2.21$), and the estimated preference parameter ($\hat{h} = 0.947$) implies that agent place high weight on relative consumption .[15]

The estimated share of project surplus that goes to innovating households is $\hat{\eta} = 0.785$, so that incomplete risk sharing is an important feature of the estimated model. The volatility of the two technology shocks is $\hat{\sigma}_x = 7.7\%$ and $\hat{\sigma}_\xi = 13.7\%$. Last, the estimate for the adjustment-cost parameter is $\hat{\alpha} = 0.36$ implying adjustment costs that are not far from quadratic at the project level. The estimated parameters governing the evolution of $\lambda_{f,t}$ imply that the high- and low- growth states have very different project acquisition rates ($\hat{\lambda}_H = 8.6, \hat{\lambda}_L = 0.12$), with the high growth state being highly transitory ($\hat{\mu}_H = 0.015, \hat{\mu}_L = 0.283$).

Not all of the parameters are precisely estimated. Their precision reflects the degree to which the output of the model is sensitive to the individual parameter values. For instance, the rate of capital depreciation, the mean values of the two technology shocks and the preference parameter θ are estimated with large standard errors. As is typically the case, shock volatilities are fairly precisely estimated.

Relative to existing general equilibrium production-based models of asset prices, the main success of our model is in replicating the cross-sectional patterns in asset returns.[16] These patterns arise

[15]Preference parameters θ and γ are high relative to the range of values typically considered in the literature. However, relative consumption preferences alter the meaning of these parameters. In particular, without randomness, the elasticity of intertemporal substitution does not equal θ. For instance, estimating the EIS via a Hall-style regression in simulated data from the model yields a mean estimate of 1.08 across simulations. Regarding the other two preference parameters, in Section 3.3.2 we show that their values can be lowered significantly by extending the model, for instance, by introducing limited asset market participation.

[16]Recent successes include Boldrin et al. (2001); Guvenen (2009); Campanale et al. (2010); Kaltenbrunner and Lochstoer (2010); Garleanu et al. (2012); Croce (2014). Some of these models succeed in generating smooth consumption paths with low-frequency fluctuations and volatile asset returns, sometimes at the cost of a volatile risk-free rate. The models are often hampered by the fact that, in response to a positive technology shock consumption rises while dividends typically fall, implying that the aggregate payout of the corporate sector is negatively correlated with

primarily from value and growth firms having differential exposures to the two technology shocks x and ξ, imperfect risk sharing, and preferences over relative consumption. We next detail the specific mechanisms that lead to these results.

3 Examining the model's mechanism

To obtain some intuition about the asset pricing predictions of the model, it is helpful to analyze the relation between technological progress, the stochastic discount factor, and asset returns. We begin our analysis in Section 3.1, where we examine how technology shocks enter into the investors' stochastic discount factor (SDF). In Section 3.2 we examine how these two shocks impact asset returns in the cross-section. In Section 3.3 we examine which of the model's non-standard features are important for the model's quantitative performance. In addition, we introduce an extension of the model that allows for limited participation in financial markets.

3.1 The pricing of technology risk

To understand how the two technology shocks x and ξ affect the equilibrium SDF, we need to establish their impact on the consumption of individual agents. Because of imperfect risk sharing, there is a distinction in how these shocks affect aggregate quantities and how they affect individual households. To emphasize this distinction, we begin our analysis by examining the impact of technology on aggregate economic growth, and continue by examining its impact on the distribution of consumption for individual households.

3.1.1 Aggregate quantities and asset prices

We compute impulse responses for aggregate output Y_t, consumption C_t, investment I_t, labor income w_t and aggregate payout to shareholders D_t to the two technology shocks x and ξ. The latter is equal to total firm profits minus investment expenditures and payout to new inventors, $D_t = p_{Z,t} Z_t - I_t - \eta \, \lambda \, \nu_t$. Payout could be negative; using the baseline parameter estimates this is almost never the case. We compute impulse responses taking into account nonlinear dynamics of the economy. The shape of these impulse responses potentially depends on the current state vector X. In our model, the scalar state variable ω summarizes all relevant information. We therefore compute impulse responses at the mean of the stationary distribution of ω.

consumption (see e.g., Rouwenhorst, 1995; Kaltenbrunner and Lochstoer, 2010). In our setup, consumption and dividends are positively correlated, which helps the model deliver a sizeable equity premium.

We plot these impulse responses in Figure 2. Panel A shows that a positive disembodied technology shock x leads to an increase in output, consumption, investment, payout and dividends. The increase in investment leads to higher capital accumulation, so the increase in output is persistent. However, most of this increase occurs on impact. Since x is complementary to existing capital, most of its benefits are immediately realized. In panel B, we plot the response of these equilibrium quantities to a technology shock ξ that is embodied in new capital. In contrast to x, the technology shock ξ affects output only through the formation of new capital stock. Consequently, it has no immediate effect on output, and only leads to a reallocation of resources from consumption to investment on impact. Further, shareholder payout declines immediately after the shock, as firms cut dividends to fund new investments. In the medium run, the increase in investment leads to a gradual increase in output, consumption, payout and the equilibrium wage (labor income).

Next, we examine the impact of technology on the pricing of financial assets and human capital – the households total marketable wealth.The total wealth of all existing households,

$$W_t \equiv \int_0^1 W_{n,t}\, dn = V_t + G_t + H_t, \tag{23}$$

equals the sum of three components. The first part is the value of a claim on the output of all *current* projects $\mathcal{J}_t$,

$$V_t \equiv \int_{j\in\mathcal{J}_t} \left(\mathrm{E}_t \left[\int_t^\infty \frac{\pi_s}{\pi_t} \left(p_{Z,s}\, \zeta_{j,s}\, ds \right) \right] \right) dj. \tag{24}$$

The second component is the value of new growth opportunities that accrues to shareholders,

$$G_t \equiv (1-\eta) \int_0^1 \left(\mathrm{E}_t \left[\int_t^\infty \frac{\pi_s}{\pi_t} \left(\lambda_{fs} \nu_s \right) ds \right] \right) df \tag{25}$$

where

$$\nu_t \equiv \max_{k_{j,t}} \mathrm{E}_t \left[\int_t^\infty \frac{\pi_s}{\pi_t}\, p_{Z,s}\, \zeta_{j,s}\, ds \right] - k_{j,t}^{1/\alpha} \tag{26}$$

denotes the value of a new project implemented at time t. The last part denotes the value of the labor services,

$$H_t \equiv \mathrm{E}_t \left[\int_t^\infty e^{-\delta^h (s-t)} \frac{\pi_s}{\pi_t}\, \mathrm{w}_s\, ds \right]. \tag{27}$$

In Figure 3 we see how technology shocks impacts the risk-free rate, the value of installed assets V_t, growth opportunities G_t, and human capital. In our subsequent analysis, the ratio of the value of new technologies ν_t to total wealth W_t plays a key role. Thus, we also plot the response of ν_t/W_t

to a positive technology shock. A positive technology shock increases expected consumption growth; hence, as we see in the first column, the risk-free rate rises on impact. The next three columns plot the response of aggregate assets in place (V_t), the present value of growth opportunities (G_t), and human capital (H_t) to a technology shock. A positive disembodied shock x is complementary to installed capital, hence the value of assets in place and growth opportunities rises on impact. By contrast, the technology shock ξ that is embodied in new capital lowers the value of existing assets V but it increases the value of growth opportunities G.

Importantly, the value of new projects ν relative to total wealth W rises in response to either technology shock, but quantitatively, the effect is much larger for the technological shock that is embodied in *new* projects compared to advances in technology that affect both existing and new projects. This difference is important because the responses of aggregate consumption in Figure 2 to technology shocks mask substantial heterogeneity. As we show next, larger changes in ν/W lead to greater reallocation of wealth among the households.

3.1.2 Technology and individual consumption

The consumption of an individual household I differs from aggregate consumption C_t due to to imperfect risk sharing. The current state of a household can be summarized by its current share of total wealth, $w_{n,t} \equiv W_{n,t}/W_t$. The functional form of preferences (12-13) together with our assumption that the scale of the household-level innovation process is proportional to individual wealth imply that optimal individual consumption and portfolio plans are proportional to individual wealth. Then, a household's consumption share is the same as its wealth share, $c_{n,t} = w_{n,t}$, and therefore individual consumption satisfies

$$C_{n,t} = C_t \, w_{n,t}. \tag{28}$$

The dynamic evolution of households' relative wealth,

$$\frac{dw_{n,t}}{w_{n,t}} = \delta^h \, dt + \frac{\lambda}{\mu_I} \frac{\eta \, \nu_t}{W_t} \left(dN_{it}^I - \mu_I \, dt \right), \tag{29}$$

where $N_{n,t}^I$ is a Poisson process that counts the number of times that household n has acquired a new blueprint.

The first term in (29) captures the flow payoff of the annuity, as is standard in perpetual youth OLG models (Blanchard, 1985). The second term capture changes in the household's wealth resulting

from innovation. Both the drift and the return to successful innovation depend on the fraction of shareholder wealth that accrues to all successful inventors, $\eta\,\nu_t/W_t$. Each period, a household yields a fraction $\lambda\,\eta\,\nu_t/W_t$ of its wealth share to successful innovators. This wealth reallocation occurs because households own shares in all firms, and these firms make payments to new inventors in return for their blueprints. During each infinitesimal time period, with probability $\mu_I\,dt$, the household is itself one of the innovators, in which case it receives a payoff proportional to $\eta\nu_t W_{n,t}$. We can thus see how an increase in the share of new technologies ν_t/W_t leads to an increase in the household's idiosyncratic risk, similar to the model of Constantinides and Duffie (1996).

Examining (29), we see that the effect of technological progress on the cross-sectional distribution of household wealth is a function of the ratio of the value of new projects to the total shareholder wealth, ν_t/W_t. Innovation reallocates wealth shares from many households to a select few. Although an increase in ν_t/W_t does not affect the *expected* wealth share of any household, it raises the magnitude of unexpected changes in households' wealth shares. Since households are risk averse, they dislike the variability of changes in their wealth (29). Preferences over relative consumption ($h > 0$) magnify the negative welfare effect of relative wealth shocks.

We next explore the impact of technology shocks on the consumption path of an individual household. Our objects of interest are the household's relative wealth share w_i, the household's consumption C_i, and household consumption adjusted for relative preferences $C_i^{1-h}w_i^h$. In the first three columns of Figure 4, we plot the impulse response of these variables to the two technology shocks. In addition to the response of the mean, we also plot how the median of the future distribution of these variables changes in response to the two technology shocks.

The first column of Figure 4 summarizes the role of incomplete risk sharing in our model. A technology shock – either x_t or ξ_t – has no impact on the expected future wealth share w_T at any horizon, because in our model technology shocks have a symmetric ex ante effect on all households. However, the lack of the average effect on the wealth share masks substantial heterogeneity in individual outcomes. Specifically, the response of the median of the distribution is significantly negative at all horizons. The very different responses of the mean and the median wealth share suggest a highly skewed effect of technological shocks on individual households, which is key to understanding the effects of technology shocks on the stochastic discount factor in our model.

As we show above, an increase in either x or ξ implies that the equilibrium value of new projects ν – scaled by total wealth W – increases. Equation (29) shows that this increase in the returns to successful innovation implies that most households experience higher rates of relative wealth decline – a reduction in the drift of dw_t – whereas the few households that innovate increase their wealth

shares greatly. From the perspective of a household at time t, the distribution of future consumption becomes more variable and more skewed following a positive technological shock, even though on average the effect is zero – both the extremely high realizations of w and the paths along which w declines persistently are now more likely. Unlike the mean, the median of w_T is not influenced by the rare but extremely positive outcomes, and instead reflects the higher likelihood of large gradual wealth declines.

The next two columns of Figure 4 examine the response of consumption. In the second column, we see that the response of the mean and median future consumption to a disembodied shock x is not that different; the difference is larger for the embodied shock – since it has a larger impact on ν/W. The next column shows the role of relative consumption preferences. If households care about their relative consumption w_i in addition to their own consumption C_i, then the impact of technology shocks on their adjusted consumption flow is a weighted average of the first two columns.

The difference between the response of the mean and the median of the distribution of future consumption highlights the asymmetric benefits of technology shocks. This difference is most stark when technology is embodied in new vintages – since it leads to larger movements in the share of value due to new technologies and thus higher returns to innovation. Since households are risk averse, the mean response is insufficient to characterize the impact of technology on their indirect utility. When evaluating their future utility, households place little weight on the extremely high paths of w. Hence the median response becomes more informative. Put differently, in addition to their effect on the mean consumption growth, technology shocks also affect the variability of consumption because they affect the magnitude of the jump term in (29). Even though the conditional risk of individual wealth shares is idiosyncratic, the conditional moments depend on the aggregate state of the economy, and therefore innovation risk affects the stochastic discount factor, just like in the model of Constantinides and Duffie (1996). To illustrate this connection, the last column of Figure 4 plots the increase in the variance of instantaneous consumption growth. We see that both technology shocks lead to higher consumption volatility. The effect is substantially higher for ξ than for x, again due to its higher impact on the returns to innovation ν/W.

In sum, Figure 4 shows that the economic growth that results from technological improvements are not shared equally across households. This is true for both technology shocks, but the effect is quantitatively larger for improvements in technology that are embodied in new capital goods. This asymmetric distribution of benefits – coupled with preferences over relative consumption – will have important implications about the pricing of these shocks, as we see next.

3.1.3 The Stochastic Discount Factor

Next, we examine the stochastic discount factor. Financial markets in our model are incomplete, since some of the shocks – the acquisition of blueprints by individual households – are not spanned by the set of traded financial assets. As a result, there does not exist a unique stochastic discount factor (SDF) in our model – as in Constantinides and Duffie (1996), the utility gradients of various agents are not identical and each can serve as a valid SDF. To facilitate the discussion of the aggregate prices, we construct an SDF that is adapted to the market filtration $\mathcal{F}$ generated by the aggregate productivity shocks (B_x, B_ξ). This SDF is a projection of agent-specific SDFs (utility gradients) on $\mathcal{F}$. The following proposition illustrates how to construct a valid SDF in our economy.

Proposition 1 (Stochastic Discount Factor) *The process π_t, given by the following equation, is adapted to the market filtration $\mathcal{F}$ and is a valid SDF:*

$$\log \pi_t = \int_0^t b(\omega_s)\, ds - \gamma_1 \chi_t - \frac{1}{\theta_1}\left(\log C_t - \chi_t\right) - \frac{1-\kappa}{\kappa} \log f(\omega_t). \tag{30}$$

In the above equation, $\kappa \equiv \frac{1-\gamma}{1-\theta^{-1}}$, $\gamma_1 \equiv 1-(1-\gamma)(1-h)$, and $\theta_1 \equiv \left(1-(1-\theta^{-1})(1-h)\right)^{-1}$. In the first term, the function $b(\omega)$ is defined in the proof of the proposition in the Appendix. In the last term, the function $f(\omega)$ is related to the value function J of an investor with relative wealth w_{it},

$$f(\omega_t) = (1-\gamma)\, J(w_{i,t}, \chi_t, \omega_t) \left(w_{i,t}^{1-\gamma} e^{(1-\gamma_1)\chi_t}\right)^{-1}. \tag{31}$$

Equilibrium prices of the aggregate technological shocks stem from four stochastic terms in the expression for the SDF in (30). We construct impulse responses for the log SDF – and its components – taking into account the nonlinear nature of equilibrium dynamics; we introduce an additional one-standard deviation shock at time t without altering the realizations of all future shocks. We plot these responses in Figure 5.

The first two components in (30) capture the permanent component of the SDF. The second term $-\gamma_1 \chi_t$ is a non-stationary process that loads on both innovation shocks with constant coefficients. The coefficients γ_1 and θ_1 are essentially the risk aversion and EIS coefficients modified towards one to account for relative consumption considerations. That is, the household prices shocks that do not affect its relative consumption w_i using γ_1 and θ_1. Recall that the term χ_t captures the permanent impact of technology shocks on the aggregate consumption process. The third term in (30) reflects the pricing of transitory shocks to aggregate consumption. This transitory deviation of

log consumption from its stochastic trend χ_t is a function of the stationary process ω_t. As we can see from the fourth column of Figure 5, the contribution of this term to the dynamics of the SDF is quantitatively small.

Most importantly, the last term in (30) captures fluctuations in investors' indirect utility—holding fixed the investor's wealth share w and stochastic trend z—resulting from technology shocks. This term summarizes the effect of technology shocks on the distribution of future consumption growth. Importantly, this term also captures the effect of incomplete markets – household expectations about the distribution of future wealth shares w_i. The impact of technology on the indirect utility function is further exacerbated by the fact that individual innovation shocks have a highly positively skewed distribution. A positive technology shock implies a higher payoff to individual households from successful innovation and a larger decline in relative consumption absent thereof. An investor with concave preferences does not place as much weight on the low-probability high-payoff outcomes as he does on the highly likely persistently low consumption growth. Thus, a positive technology shock can have a large negative effect on the indirect utility of the investors. As we can see from the impulse responses shown in Figure 5, this is a quantitatively important aspect of how technology shocks are priced in our model – in particular for the embodied shock ξ. This difference in magnitude in the response of indirect utility arises because the reallocative effects of an embodied shock – its impact on the returns to innovation ν_t/W_t – is an order of magnitude higher than the reallocative effects of the disembodied shock, as we saw in the last column of Figure 3.

Comparing panels A and B of Figure 5, we see that the two technology shocks carry opposite prices of risk in our model. A positive disembodied shock x negatively affects the SDF on impact, implying that it carries a positive risk premium. By contrast, a positive embodied shock ξ leads to a rise in the SDF on impact, implying a negative risk premium. As a result, households value securities that provide a hedge against states of the world when ξ is high and x is low. This difference in how the SDF responds to the two technology shocks stems primarily from the response of the indirect utility term $f(\omega_t)$ in the SDF. Both technology shocks x and ξ lead to an increase in the permanent component χ_t of consumption, which by itself causes the SDF to fall. However, in the case of the embodied shock, the fall in indirect utility due to the unequal sharing of benefits from technological progress is sufficiently large to offset the benefits of higher aggregate consumption. The resulting demand for insurance against high realizations of ξ is driven by the endogenous increase in the consumption uncertainty of individual investors.

3.2 Technology shocks and the cross-section of asset returns

Next, we examine the impact of technology shocks on individual firms. The firm's current state is fully characterized by the aggregate state X_t, its probability of acquiring new projects $\lambda_{f,t}$, its relative size,

$$k_{f,t} \equiv \frac{K_{f,t}}{K_t} = \frac{1}{K_t} \sum_{j \in \mathcal{J}_{f,t}} e^{\xi_{s(j)}} k_{j,t}. \tag{32}$$

and its current average productivity of projects

$$\bar{u}_{f,t} \equiv \frac{Z_{f,t}}{K_{f,t}} = \left(\sum_{j \in \mathcal{J}_{f,t}} e^{\xi_{s(j)}} \, u_{j,t} \, k_{j,t} \right) \Bigg/ \left(\sum_{j \in \mathcal{J}_{f,t}} e^{\xi_{s(j)}} k_{j,t} \right). \tag{33}$$

Next, we examine how the impact of technology shocks on firm outcomes varies with their current state, $(\lambda_{ft}, k_{ft}, \bar{u}_{ft})$.

3.2.1 Technology and creative destruction

In our model, technological progress indirectly leads to displacement of installed capital due to general equilibrium effects. To see this, consider the effect of technological progress on firm profitability. Competition across firms implies that a firm's revenue flow,

$$p_{Z,t} Z_{f,t} = \phi \, Y_t \, z_{f,t}, \quad \text{where} \quad z_{ft} \equiv Z_{f,t}/Z_t = \bar{u}_{f,t} \, k_{f,t}. \tag{34}$$

is a function of the aggregate state X_t, and of the firm's current revenue share $z_{f,t}$ – which in the absence of ongoing costs of production also equals its share of total profits. The firm's market share evolves according to

$$\frac{dz_{f,t}}{z_{f,t}} = \kappa_u \left(\frac{1}{\bar{u}_{f,t}} - 1 \right) dt + \lambda \, a_0 \, \frac{\nu_t}{V_t} \left(\frac{\lambda_{f,t}}{\lambda \, z_{f,t}} - 1 \right) dt + a_0 \, \frac{\nu_t}{V_t} \left(dN_{f,t} - \lambda_{f,t} \, dt \right) + \sigma_u \, dB^u_{f,t}, \tag{35}$$

where a_0 is a constant.

The first term in equation (35) captures the effect of mean-reversion in profitability – project-specific shocks u have mean equal to 1. The second term in (35) captures fluctuations in expected share of output due to changes in technology. A positive embodied technology shock ξ, or equivalently an increase in investment I, raises the output share for firms with high growth potential (λ_{ft}) but low current output (Z_{ft}). These are the firms that, in comparison to an average firm in the economy, derive most of their value from their future growth opportunities rather than their existing

operations. We refer to such firms as growth firms. Conversely, firms that have a large scale of operations relative to their future growth potential – value firms – will on average experience a decline in their output share – and hence profitability – following a positive shock to ξ. The last two terms in (35) capture idiosyncratic shocks to firm profitability.

To illustrate the resulting difference in firm outcomes, we examine the impact of the two technology shocks x and ξ on two firms with high and low levels of λ_f/z_f, which we term growth and value firms, respectively. To isolate the effect of mean-reversion in firm productivity, we set $\bar{u}_f = 1$ for both firms. Examining Figure 6 , we note a stark difference in firm outcome.

In sum, growth firms have higher stock return exposure to either technology shock than value firms – the difference being quantitatively larger for embodied versus disembodied shocks. As we see in panel A, improvements in technology that are complementary to all capital lead to an increase in profitability for both firms on impact. The growth firm is more likely to have higher investment opportunities than the value firm, hence on average they increase investment and pay lower dividends in the short run. In terms of stock price reactions, growth firms experience higher returns than value firms. In panel B, we see that value and growth firms have very different responses to technology improvements embodied in new vintages (ξ). The technology shock ξ leaves the output of existing projects unaffected – it only increases the productivity of new investments. Due to the equilibrium response of the price of output p_Z, cashflows from existing operations fall. Growth firms increase investment, and experience an increase in profits in market valuations. Value firms have few new projects to invest in, hence their profits and valuations fall.

3.2.2 Firm risk exposures and the market-to-book ratio

We next examine how these differential responses to technology shocks into cross-sectional differences in risk and risk premia. A major challenge in examining the predictions of the model for the cross-section of firms is that the firm's current state $(\lambda_{ft}, k_{ft}, \bar{u}_{ft})$ is unobservable. Hence, we have to rely on observable proxies. The most commonly used measure of a value or growth firm is its ratio of market value of the firm to its book value of assets K_{ft}.[17] In a world without financial leverage, this ratio is also equal to the firm's Tobin's Q. In our model, a firm's log market to book ratio can

[17] In our model, the replacement cost of capital – the equivalent amount of new capital that yields the same output – is $e^{-\xi_t} K_{ft}$. Since we are focusing on cross-sectional differences, the exact price deflator applied to the book value of assets is immaterial.

be written as

$$\log Q_{ft} - \log Q_t = \log \left[\frac{V_t}{V_t + G_t} \left(1 + \widetilde{p}(\omega_t)\ (\bar{u}_{ft} - 1)\right) + \frac{G_t}{P_t + G_t} \frac{\bar{u}_{ft}}{z_{ft}} \left(1 + \widetilde{g}(\omega_t)\ \left(\frac{\lambda_{ft}}{\lambda} - 1\right)\right)\right], \tag{36}$$

where Q_t is the market-to-book ratio of the market portfolio and V_t and G_t, defined in equations (24) and (25) respectively, are functions only of the aggregate state.

Examining (36), we note that a firm's market to book ratio is increasing in the likelihood of future growth λ_f, decreasing in the firm's relative size z_f, and increasing in the firm's average productivity $\bar{u}_f$. This latter effect prevents Tobin's Q from being an ideal measure of growth opportunities, since it is contaminated with the profitability of existing assets. We next examine the extent to which Q is a useful summary statistic for firm risk and risk premia in our model. To do so, we examine how changes in the firm's current state $(\lambda_{ft}, k_{ft}, \bar{u}_{ft})$ jointly affect both firm Q and risk exposures. We plot the results in Figure 7. In each of the three columns, we vary one of the elements of the firm's current state (λ_{ft}, z_{ft}, or $\bar{u}_{ft}$) and keep the other two constant at its steady-state mean. On the horizontal axis, we plot the change in the firm's Q (relative to the market). On the vertical axis we plot the firm's return exposure to x and ξ (panels A and B, respectively) and the firm's risk premium (panel C). We scale the horizontal axis so that it covers the 0.5% and 99.5% of the steady-state distribution of the each of the firm's state variables. The resulting cross-sectional distribution of Tobin's Q is highly skewed.

Examining Figure 7, we see that regardless of the source of the cross-sectional dispersion in Q, the relation between Q and risk exposures is positive. The pattern in the first two columns is consistent with the impulse responses in the previous section – small firms with high probability of acquiring future projects have higher technology risk exposures than large firms with low growth potential. The last column shows that increasing the firm's average profitability $\bar{u}_f$ – holding λ_f and k_f constant increases both Q and risk exposures.[18] However, the magnitude of this effect is quantitatively minor.

The last row of Figure 7 shows how the firms' risk premium (unlevered) is related to cross-

[18]This pattern might seem puzzling initially in light of the discussion in the previous section, since increasing productivity $\bar{u}_f$ while holding size k_f and investment opportunities λ_f constant will lower λ_f/z_f. However, altering $\bar{u}_f$ also has a cashflow duration effect: due to mean reversion in profitability, profitable firms have lower cashflow duration – their cashflows are expected to mean-revert to a lower level. This lower duration of high u_f firms implies that their valuations are less sensitive to the rise in discount rate following a positive technology shock– see in Figure 3. In our calibration, this duration effect overcomes the effect due to λ_f/z_f, implying a somewhat more positive stock price response for high $\bar{u}_f$ firms.

sectional differences in Q. Recall that the two technology shocks carry risk premia of the opposite sign. The disembodied shock carries a positive risk premium; in the absence of other technology shocks, it would would imply that firm risk premia rise with market-to-book. However, the fact that the embodied shock carries a negative risk premium – coupled with its higher volatility – imply a lower risk premium for growth firms relative to value firms. Households are willing to accept lower average returns for investing in growth firms because doing so allows them to partially hedge the displacement arising from the embodied shock ξ – the decline in their continuation utility.

3.3 Sensitivity analysis and extensions

Here, we explore the impact of alternative formulations of the model on our results. In particular, our model has three relatively non-standard features on the household side. First, our model features technology shocks that are embodied in new capital. Second, markets are incomplete in that households cannot sell claims on their proceeds from innovation. Third, household preferences are affected by their consumption relative to the aggregate economy. Here, we examine how important these features are for the quantitative performance of the model. In addition, we estimate an extended version of the model that allows for limited stock market participation.

3.3.1 Sensitivity to modelling assumptions

We estimate three restricted versions of the model. The first version assumes complete markets over innovation outcomes, or equivalently $\eta = 0$. In this case, all proceeds from new projects accrue financial market participants. The second restricted model constrains $h = 0$ so that households have no preferences over relative consumption. The third restricted model features no embodied technological change – we restrict $\mu_\xi = 0$ and $\sigma_\xi = 0$. To estimate these models, we repeat the procedure detailed in Section 2. We report the performance in matching the target set of empirical statistics of the restricted versions in columns (R1) through (R3) of Table 3. We report the corresponding parameter estimates in column (R1)-(R3) of Table 4.

We see that all three of these assumptions play an important role. The model with complete markets (R1) generates essentially zero risk premia, both in average (the market portfolio) and also in the cross-section (the value factor). The model without relative preferences (R2) generates realistic aggregate dynamics and moments of the market portfolio, and in contrast to models (R1) and (R3) uses only moderate levels of utility curvature ($\hat{\gamma} = 17.8$). However, (R2) generates virtually no cross-sectional dispersion in risk premia. Last, the model with only disembodied technology

shocks (R3) has difficulty jointly matching the dynamics of aggregate consumption and investment. It does generate a sizable equity premium, however it generates almost no cross-sectional dispersion in risk premia between value and growth firms.

3.3.2 Limited participation

Our baseline model is quite successful in matching empirical facts about economic quantities and asset returns. However, some of the parameter estimates may seem implausibly high based on introspection – for instance, the high degree of risk aversion, the preference weight on relative consumption and the share of the surplus that accrues to innovators.

Here, we illustrate an additional channel through which technology shocks can lead to displacement of financial market participants. As we saw in Figure 2, labor income rises in response to improvements in technology; this rise in labor income acts as a natural hedge for the displacement of households that occurs through financial markets. The fact that labor income acts as a hedge is an artifact of the stylized nature of our model – technology has no displacive effect on labor. A more realistic model that allows for endogenous displacement of human capital is outside the scope of this paper. However, recognizing that only a subset of households participates in financial markets serves to mitigate this channel. For instance, Poterba and Samwick (1995) report that the households in the top 20% have consistently owned more than 98% of all stocks. If most workers do not participate in financial markets, the share of labor income accruing to market participants can be small. We thus estimate an extension of our baseline model that allows for limited participation in the stock market.

We model limited participation by assuming that newly born households are randomly assigned to one of two types, *shareholders* (with probability q_S) and *workers* (with probability $1 - q_S$). Shareholders have access to financial markets, and optimize their life-time utility of consumption, just like the households in our baseline model. Workers in this economy are hand-to-mouth consumers. They do not participate in financial markets, supply labor inelastically and consume their labor income as it arrives. Workers can also successfully innovate (just like shareholders); those that do so become shareholders. To conserve space, we relegate the details of this extension to the Appendix. To estimate the extended model, we include two additional target statistics that help identify q_S. First, we target a mean consumption share of 42.9% based on the estimates of Guvenen (2006). Second, we target a volatility of shareholder consumption growth of 3.7% based on the unpublished working paper version of Malloy, Moskowitz, and Vissing-Jorgensen (2009) –

which includes an adjustment for measurement error.

We report the performance of the extended model in matching the target set of empirical statistics of the restricted versions in column (X1) of Table 3 and the corresponding parameter estimates in column (X1) of Table 4. We see that the extended model does indeed do better in matching the empirical facts with lower levels of risk aversion ($\hat{\gamma} = 58$ vs 106), share of surplus to innovators ($\hat{\eta} = 0.6$ vs 0.79) and preference share over relative consumption ($\hat{h} = 0.8$ vs 0.95). The estimate $\hat{q}_s = 0.073$ implies that the steady-state fraction of households that participate in financial markets is approximately 12%, consistent with the evidence in Guvenen (2006), who reports that the top 10% of households own approximately 84% of all financial assets.

3.3.3 Discussion

The fact that our assumption about relative consumption preferences plays an important role deserves some justification. From a pure modeling standpoint, this assumption allows for differential pricing of shocks that affect all agents symmetrically, versus reallocative shocks that lead to cross-sectional dispersion in household consumption. Recalling proposition 1, the former (level) shocks are priced according to the modified preference parameters $[\gamma_1, \theta_1]$, whereas the reallocative shocks are priced using $[\gamma, \theta]$. Relative consumption preferences allows the model to price these two shocks differently in order to match the observed features of asset returns. To this end, we see that most estimated versions of the model require high values of γ, but since the estimate of h is typically high, the implied coefficient γ_1 is much smaller. We speculate that an alternative mechanism to allow agents to price level and reallocative shocks differently is to have preferences that feature ambiguity aversion, and allow the level of ambiguity surrounding the level versus reallocative effects of technology to differ. This is an interesting avenue for future research that however lies outside the scope of the current paper.

4 Additional Predictions

Here, we examine the performance of the model in replicating some features of the data that we do not use as an explicit estimation target in Section 2. In Section 4.1, we consider the correlation between consumption growth and asset returns between the model and the data. In Section 4.2, we develop additional direct tests of the model mechanism. To conserve space, we describe our tests briefly and provide further details in the Appendix.

4.1 Consumption and asset returns

A natural way to test the model is to examine its implications for the joint distribution of consumption and asset returns. This joint distribution is not among our estimation targets. In addition to the results for the baseline model, we also report the results for the model with limited market participation in Section 3.3.2. Motivated by that extension, we examine the correlation of asset returns with the consumption of households that participate in the stock market. We use the data of (Malloy et al., 2009) which is available over the 1982-2002 period. We compute correlations of asset returns and the consumption of stockholders in absolute terms, but also relative to the consumption of non-participants. The behavior of the latter variable plays an important role in the model with limited participation. We follow standard practice and aggregate consumption growth over multiple horizons. We report results using 2-year growth rates, but the results are qualitatively similar using longer horizons.

Table 5 shows that the baseline model generates empirically plausible levels of correlation between dividends and aggregate consumption. Consumption and aggregate stock market returns are more highly correlated in the model, but the difference between the correlations in the model and in the data is not statistically significant. The model also reproduces the low empirical correlation between aggregate consumption and the value factor, which contributes to the failure of the Consumption CAPM to capture the value premium in simulated data. Further, we see that the extended model with limited market participation replicates one of the key findings of (Malloy et al., 2009) – returns on value firms covary more with the consumption of stockholders than returns on growth firms. However, the extended model does fail to capture the relatively low correlation of shareholder consumption and the market portfolio (equal to 21% in the 1982-2002 sample versus 72% in the model).

4.2 Direct tests of the model mechanism

In our analysis so far, we have been following closely the existing literature and evaluating the success of the model based of the model-implied correlations between the macroeconomic quantities and prices. Part of the model mechanism relies on fluctuations in the value of new projects ν_t – a quantity that is challenging to observe. Here we construct an empirical measure for the value of new technologies based on prior research. We then use this empirical proxy to directly test the model mechanism.

4.2.1 Estimating the value of new technologies

The market value of new projects ν_t plays a key role in the model's predictions, both for the dynamics of firm values (35) and for the evolution of investors' wealth (A.42). To test the model's mechanism, we use data on patents and stock returns to construct an empirical proxy for ν_t. While not all technological improvements get patented, our empirical tests confirm that the time series of the patent-based measure is informative about the overall pace of technological innovation in the economy.

We follow Kogan et al. (2012) and estimate the net present value of a patent as the change in the dollar value of the firm around a three-day window after the market learns that the firm's patent application has been successful.[19] To replicate this construction in simulated data, we employ an approximation that does not require estimation of new parameters.[20] Using this approximate construction, we can replicate their main empirical results.[21] We report one such result here. Kogan et al. (2012) document a strong positive relation between estimated patent values and future citations the patent receives–a commonly used indicator of the scientific value of new patents. To evaluate the quality of our approximation, we replicate their result using our approximation $\hat{\nu}_j$ and plot the results in panel A of Figure 8. Comparing panel A to the corresponding figure in their paper, we find a qualitatively similar relation between $\hat{\nu}_j$ and future citations.

4.2.2 Technological innovation and firm cashflows

We begin by studying the model's predictions about the cross-section of firms. As we show in equations (34-35), the impact of technology shocks on firms' cash flows is a function of the ratio ν_t/V_t, the aggregate value of new projects scaled by the value of installed capital. Kogan et al. (2012) document substantial heterogeneity in innovation across industries. Here, we exploit this heterogeneity to get sharper estimates of the impact of technological progress on firm outcomes.

[19] The dollar reaction around the issue date is an understatement of the dollar value of a patent. The market value of the firm is expected to changes by an amount equal to the NPV of the patent times the probability that the patent application is unsuccessful. This probability is not small; in the data less than half of the patent applications are successful. See Kogan et al. (2012) for an extensive discussion of the empirical procedure.

[20] Kogan et al. (2012) allow for movements in stock returns around the announcement window that is unrelated to the value of the patent. They construct a filter of the estimated patent value using specific distributional assumptions, and propose a methodology to empirically estimate those parameters using high-frequency data. This is quite difficult to replicate in simulated data because it requires re-estimating parameters in each sample. Instead, we recognize that their optimal filter can be approximated by the function $f(x) = max(x, 0)$.

[21] The details are available upon request, and also appear in earlier working paper versions of Kogan et al. (2012).

Specifically, we construct a direct analogue of ν/V at the industry level as

$$\hat{\omega}_{I\setminus f} = \frac{\sum_{f'\in I\setminus f}\sum_{j\in P_{f't}}\hat{\nu}_j}{\sum_{f'\in I\setminus f} M_{f't}}, \tag{37}$$

where industry I is defined by the 3-digit SIC code and P_{ft} is the set of patents issued to firm f in year t. $I \setminus f$ denotes the set of all firms in industry I excluding firm f. Since the value of assets is unobservable, we scale by total stock market capitalization M instead . In the model, ν_t/V_t and $\nu_t/(V_t + G_t)$ (as well as ν_t/W_t) are monotonically increasing functions of the same state variable ω_t.

We are mainly interested in how the relation between technology shocks and future firm profitability depends on the current state of the firm, that is, the interaction of ν_t/V_t with λ_f/z_f in (35). We estimate the impact of technology on log firm profitability s_{ft} using the following approximation to (35),

$$s_{ft+T} - s_{ft} = (a_0 + a_1\, q_{ft})\, \hat{\omega}_{I\setminus ft} + a_t + a_I + c_1\, G_{ft} + c_2\, \hat{N}_{ft} + c_l\, s_{ft} + c_k\, k_{ft} + \varepsilon_{t+T}. \tag{38}$$

Based on the discussion in Section 3.2.2, we allow this relation to vary as a function of the firm's Tobin's Q. We classify firms as either value ($q_{ft} = 0$) or growth ($q_{ft} = 1$) depending on whether their Tobin's Q falls below or above the industry median at time t. We construct firm innovation outcomes as, $\hat{N}_{ft} = \sum_{j\in P_{ft}} \hat{\nu}_j/M_{ft}$ where $\hat{\nu}$ is our estimate of patent value constructed in Section 4.2.1. This variable controls for innovation outcomes by firm f – that is, the analogue of the dN_{ft} term in (35). Consistent with (35) we also include controls for lagged log profits y_{ft} and log size k_{ft}.

In the empirical specification, we include time and industry dummies, and cluster the standard errors at the firm level. We also estimate equation (5) in simulated data from the model. Since the model contains no industries, we drop the time and industry dummies. We scale $\hat{\omega}_{I\setminus f}$ to unit standard deviation. Table 6 compares the estimated coefficients a_0 and a_1 across horizons of 1 to 7 years in the data and the model.

In sum, our empirical estimates support the notion that improvements in technology benefit firms that have the ability to implement these technologies at the expense of firms that do not, and the model comes close to matching the magnitude of the empirical relations. The estimated coefficient a_0 in the data (panel A) is negative and statistically different from zero, implying that the impact of technological progress on the expected profitability of low-Q firms is negative. Further, the estimated coefficient a_1 is positive and statistically significant across the horizons,

revealing substantial heterogeneity in how profits of low-Q and high-Q firms respond to an increase in innovation at the industry level. The estimates of a_0 in the model (panel B) are similar in magnitude to the empirical values. The model implies somewhat larger heterogeneity between value and growth firms – the estimated coefficient a_1 is higher in simulated data. However, there is significant variation across simulated samples, and the empirical estimates lie inside the 90% confidence intervals implied by model simulations.

4.2.3 Innovation and aggregate dynamics

We next study the relation between our estimate of technological progress, aggregate consumption, and stock returns. We construct an estimate of the average value of new technologies at time t as

$$\hat{\omega}_t = \log\left(\frac{\sum_{j\in P_t}\hat{\nu}_j}{M_t}\right), \tag{39}$$

where P_t denotes the set of patents granted to firms in our sample in year t and M_t denotes the market capitalization of all firms in our sample at time t. We plot the constructed $\hat{\omega}_t$ in Panel B of Figure 8. Similar to Kogan et al. (2012), this time-series series lines up well with the three major waves of technological innovation in the U.S. – the 1930s, 1960s and early 1970s, and 1990s and 2000s.

In Table 7, we compare the empirical properties of the innovation series $\hat{\omega}_t$ to those implied by the model. In the data, $\hat{\omega}_t$ is more volatile and less persistent than in the model, consistent with the presence of substantial measurement error. We compute correlations of $\Delta\hat{\omega}$ with the aggregate consumption growth, with the consumption growth of stock-holders (as in Section 4.1 above), and with stock market returns and with the value factor.

The correlations in the model are largely consistent with their empirical counterparts. In the data, there is a negative (-50%) correlation between innovation shocks $\Delta\hat{\omega}$ and stock market returns. The value factor and $\Delta\hat{\omega}$ have correlation of -28%, implying that value stocks covary more negatively with innovation shocks than growth stocks. The model counterparts to the above correlations are -42% and -35% respectively.

The correlation between $\Delta\hat{\omega}_t$ and aggregate consumption in the data is also negative and equal to -20%, while in the model this correlation is approximately zero. In the data, the negative relation between innovation shocks and consumption is stronger for the consumption series of shareholders (-30%) or the consumption of stockholders relative to non-stock holders (-35%). These correlations are broadly consistent with both the baseline model and the extension with limited participation.

4.2.4 Stochastic discount factor

Tests using linearized pricing models are common in the empirical asset pricing literature. Here we use the estimates of the linearized stochastic discount factor (SDF) as a reduced-form summary of the model's implications for asset prices.

We specify the linearized SDF as

$$\hat{m}_t = \Delta \log \hat{\pi}_t = a - b_1 \, \Delta\hat{\chi}_t - b_2 \, \Delta\hat{\omega}_t. \tag{40}$$

Our main coefficient of interest is b_2, which captures how the value of new technologies affects the stochastic discount factor, or the risk premium associated with $\Delta\hat{\omega}$. As proxies for the trend shock $\Delta\hat{\chi}$ we use the annual growth rate in aggregate consumption, output or total-factor productivity (adjusted for utilization) from Basu, Fernald, and Kimball (2006). We follow standard practice estimate the parameters b_1 and b_2 using the moment conditions $E[\hat{m}_t \, R_t^e] = 0$, where R^e are the excess returns on 10 portfolios of firms sorted by book-to-market. We normalize a so that $E[\hat{m}_t] = 1$, which reduces the moment conditions to $E[R_t^e] = -cov(\hat{m}_t, R_t^e)$ In addition to the estimated parameters b_1 and b_2, we also report the cross-sectional R^2, the mean absolute pricing error (MAPE), Hansen's J-test and its associated p-value. We construct empirical confidence intervals for the latter three statistics using a jackknife estimator.

Table 8 compares estimates of (40) in the data and in the model. To facilitate this comparison, we normalize $\Delta\hat{\chi}$ and $\Delta\hat{\omega}$ to unit standard deviation. Examining Table 8, we see that specifications of the SDF that include only consumption, output or TFP growth – those typically implied by the existing general equilibrium models – are rejected by the data. The same is true in our model. By contrast, the two-factor specification (40) captures return differences on the test portfolios relatively well. Moreover, the magnitude of the estimated coefficients b_w is quite similar between the model and the data.

4.2.5 Innovation and inequality

A key feature of our model is that the benefits of innovation are unequally distributed. This mechanism is summarized by Figure 4, which shows that improvements in technology have very different effects on the *average* versus the *median* future consumption path of an individual household. Here, we provide supporting evidence that demonstrates that technological improvements are indeed associated with a decline in the median consumption relative to the average.

We do so using the following specification,

$$(c_T^{md} - c_t^{md}) - (c_T - c_t) = b_0 + b_1\,(\hat{\omega}_T - \hat{\omega}_t) + b_2\,(c_t^{md} - c_t) + b_3\,\hat{\omega}_t + u_t. \tag{41}$$

Here, c_t refers to log aggregate per capita consumption. We estimate the log median per capita consumption c_t^{md} in year t using the Consumer Expenditure Survey (CEX).[22] We estimate two versions of c_t^{md}. The first is based on the log total consumption expenditures of households in year t divided by the number of members The second computes c_t^{md} as the median residual from a regression of log household consumption on household-level controls (household size, years of schooling and age of primary earner). The sample covers the 1982-2010 period. As before, we also estimate 41 using simulated data from the model – each simulation contains 10,000 households over 28 years.

We report the estimated coefficients b_1 in Table 9 for horizons of one to five years. Consistent with the model, the empirical estimates of $\hat{b}_1$ are negative. However, given the short length of the sample, the coefficients are fairly imprecisely estimated and not statistically different from zero for some horizons. The same is true in simulated data. Comparing the third row (Model) with the first two rows (Data), we see that the empirical estimates are fairly close in magnitude to those implied by the model. However, given the wide confidence intervals this test of the model has low power. We conclude that estimates of (41) are qualitatively consistent with the model but that a quantitative comparison is challenging.

5 Conclusion

We develop a general equilibrium model to study the effects of innovation on asset returns. A distinguishing feature of innovation is that its benefits are not shared symmetrically across all agents in the economy. Hence, focusing on aggregate moments obscures the effects of innovation in the cross-section of both firms and households. Specifically, technological improvements embodied in new capital benefit agents that are key in the implementation of these ideas, while displacing existing firms and their shareholders. This displacement process is uneven. Firms well-positioned to take advantage of these opportunities benefit at the expense of firms unable to do so. Existing shareholders value firms rich in growth opportunities despite their low average returns, as they provide insurance against displacement.

[22] Using the full consumption distribution of the CEX is problematic for our purposes, since the CEX tends to undersample rich households. However, the effect of this undersampling on the median should be minor.

Our model delivers rich cross-sectional implications about the effect of innovation on firms and households, and our empirical analysis provides favorable evidence. We test the model's predictions using a direct measure of technological innovation using data on patents and stock returns. We find that growth firms fare better than value firms in industries that experience a wave of new innovations. In addition, we provide evidence that innovation is associated with increases in inequality. Both of these empirical relations support the predictions of our model, both qualitatively and quantitatively.

Our work suggests several avenues for future research. Technological progress tends to disrupt traditional methods of production leading to periods of increased uncertainty. This uncertainty will impact the decisions of agents with preference for robust control in non-obvious ways. Further, certain technological advances typically require particular skills to be effective. In light of recent evidence on the polarization of the job market (Autor, Katz, and Kearney, 2006), quantifying the role of technological progress as a determinant of the risk of human capital is particularly important.

Appendix A: Proofs and Derivations

Lemma 1 (Ergodic distribution for u) *The process u, defined as*

$$du_t = \kappa_u(1-u_t)\,dt + \sigma_u\, u_t\, dB_t^u \tag{A.1}$$

has a stationary distribution given by

$$f(u) = cu^{-2-\frac{2\,\kappa_u}{\sigma_u^2}} \exp\left(-\frac{2\theta}{u\sigma_u^2}\right), \tag{A.2}$$

where c is a constant that solves $\int f(u)du = 1$. Further, as long as $2\,\kappa_u \geq \sigma_u^2$, the cross-sectional variance of u is finite.

Proof. The forward Kolmogorov equation for the transition density $f(u,t)$ for this process is given by

$$\frac{\partial}{\partial t} f(u,t) = -\kappa_u \frac{\partial}{\partial u}\left[(1-u)f(u,t)\right] + \frac{1}{2}\sigma_u^2 \frac{\partial^2}{\partial u^2}\left[u^2 f(u,t)\right] \tag{A.3}$$

Since we are interested in the steady-state distribution, we replace the transition density $f(u,t)$ with the steady state one, $f(u)$, which yields the ODE:

$$0 = -\kappa_u \frac{\partial}{\partial u}\left[(1-u)f(u)\right] + \frac{1}{2}\sigma_u^2 \frac{\partial^2}{\partial u^2}\left[u^2 f(u)\right] \tag{A.4}$$

Integrating the above with respect to u yields

$$k = -\kappa_u\left[(1-u)f(u)\right] + \frac{1}{2}\sigma_u^2 \frac{\partial}{\partial u}\left[u^2 f(u)\right] \tag{A.5}$$

where k is a constant of integration that should equal zero given the requirement that $\lim_{u\to\infty} f(u) = 0$ and $\lim_{u\to\infty} f'(u) = 0$. Define $h(u) = u^2 f(u)$, we get the equation

$$\frac{2\,\kappa_u}{\sigma_u^2}\,\frac{1-u}{u^2} = \frac{h'(u)}{h(u)} \tag{A.6}$$

Integrating wrt to u gives us

$$\log h(x) = \frac{2\,\kappa_u}{\sigma_u^2}\left(-\frac{1}{u} - \log u\right) + \log c \tag{A.7}$$

where c is a constant. This implies that

$$f(u) = cu^{-2-\frac{2\,\kappa_u}{\sigma_u^2}} \exp\left(-\frac{2\theta}{u\sigma_u^2}\right). \tag{A.8}$$

The constant c solves

$$c\int_0^\infty u^{-2-\frac{2\,\kappa_u}{\sigma_u^2}} \exp\left(-\frac{2\theta}{u\sigma_u^2}\right) du = 1. \tag{A.9}$$

The integral in (A.10) is finite as long as $\kappa_u > 0$. The last part of the proof is to show that the variance of u is finite and positive. Given the solution for c, we have that

$$\int_0^\infty (u-1)^2\, c\, u^{-2-\frac{2\,\kappa_u}{\sigma_u^2}} \exp\left(-\frac{2\theta}{u\sigma_u^2}\right) du = \frac{\sigma_u^2}{2\,\kappa_u - \sigma_u^2}. \tag{A.10}$$

which is positive and finite as long as $2\,\kappa_u - \sigma_u^2 > 0$.

■

Lemma 2 (SDF in incomplete markets) *Consider an infinitely-lived investor with initial financial wealth W_0. The market consists of I financial assets that pay no dividends. Let S_t denote the vector of prices of the financial assets. S_t is an Ito process*

$$dS_t = \mu_t\,dt + \sigma_t\,dB_t. \tag{A.11}$$

The first asset is risk-free, and earns the interest rate r_t. Let $\mathcal{F}$ denote the natural filtration generated by the Brownian motion vector B_t. The investor receives a flow of income from innovation projects according to an exogenous Poisson process N_t with the arrival rate λ. We assume that upon arrival of such a project, investor's wealth increases by a factor of $\exp(v_t)$, where the process v is adapted to the filtration $\mathcal{F}$. The investor maximizes Epstein-Zin utility characterized by equations (12)-(13) in the main text, subject to the dynamic budget constraint

$$dW_t = \delta W_t\,dt + (e^{v_t} - 1)\,W_t\,dN_t - c_t\,dt + \theta_t dS_t, \quad W_t = \theta_t S_t, \tag{A.12}$$

and a credit constraint, which rules out doubling strategies and asymptotic Ponzi schemes:

$$W_t \geq 0. \tag{A.13}$$

Let $C_t^\star$, $\theta_t^\star$, and $W_t^\star$ denote the optimal consumption strategy, portfolio policy, and the wealth process respectively. Let $J_t^\star$ denote the utility function under the optimal policy. Then π_t, defined by

$$\pi_t = \exp\left(\int_0^t \delta + \frac{\partial\phi(C^\star{}_s, J^\star{}_s; \bar{C}_s)}{\partial J^\star{}_s} + \lambda\left(e^{(1-\gamma)v_s} - 1\right)\,ds\right)\,A_t, \tag{A.14}$$

where

$$A_t = \frac{\partial\phi(C_t^\star, J_t^\star; \bar{C}_t)}{\partial C_t^\star}\exp\left(\int_0^t \gamma v_s\,dN_s\right) \tag{A.15}$$

is adapted to filtration $\mathcal{F}$ and is a state-price density process in this economy.

Proof. Let $\mathcal{M}$ denote the market under consideration, and define a fictitious market $\hat{\mathcal{M}}$ as follows. $\hat{\mathcal{M}}$ has the same information structure as $\mathcal{M}$, with modified price processes for financial assets. Specifically, let

$$G_t = \exp\left(\int_0^t \delta\,ds + v_s\,dN_s\right) \tag{A.16}$$

and define price processes in the market $\hat{\mathcal{M}}$ as

$$\hat{S}_t = G_t S_t. \tag{A.17}$$

The budget constraint in the market $\hat{\mathcal{M}}$ is standard,

$$d\hat{W}_t = -c_t\,dt + \hat{\theta}_t d\hat{S}_t, \quad \hat{W}_t = \hat{\theta}_t\hat{S}_t. \tag{A.18}$$

If a consumption process $\{c_t\}$, $t \in [0, \infty)$ can be financed by a portfolio policy θ_t in the original market $\mathcal{M}$, it can be financed by the policy $G_t^{-1}\theta_t$ in the fictitious market $\hat{\theta}_t = G_t^{-1}\theta_t$, and vice versa. Thus, the set of feasible consumption processes is the same in the two markets, and therefore the optimal consumption processes are also the same. Since the consumption-portfolio choice problem in the fictitious market is standard, the utility gradient of the agent at the optimal consumption policy defines a valid state-price density process $\hat{\pi}_t$,

$$\hat{\pi}_t = \exp\left(\int_0^t \frac{\partial\phi(C^\star{}_s, J^\star{}_s; \bar{C}_s)}{\partial J^\star{}_s}\,ds\right)\,\frac{\partial\phi(C_t^\star, J_t^\star; \bar{C}_t)}{\partial C_t^\star}. \tag{A.19}$$

Thus, for all $t < T$,

$$\hat{\pi}_t G_t S_t = \hat{\pi}_t \hat{S}_t = \mathrm{E}_t \left[\hat{\pi}_T \hat{S}_T\right] = \mathrm{E}_t \left[\hat{\pi}_T G_T S_T\right] \tag{A.20}$$

and therefore $\pi'_t = \hat{\pi}_t G_t$ is a valid state-price density in the original market $\mathcal{M}$.
Note that π'_t is not adapted to the filtration $\mathcal{F}$, since it depends on the agent's innovation process N. In other words, π'_t is an agent-specific state-price density process. The last remaining step is to show that the process π_t, which is adapted to the filtration $\mathcal{F}$, is also a valid SPD. Next, we show that the process $\exp\left(\int_0^t \partial\phi(C_t, J_t; \bar{C}_t)/\partial J_t\right)$ is adapted to $\mathcal{F}$, and the process $\partial\phi(C_t, J_t; \bar{C}_t)/\partial C_t$ can be decomposed as $A_t G_t^{-\gamma}$, where A_t is also adapted to $\mathcal{F}$. Given the homotheticity of the Epstein-Zin utility function and the budget constraint (A.12), standard arguments show that the agent's value function and the optimal consumption policy can be expressed as

$$J_t^\star = (W_t^\star)^{1-\gamma} \kappa_t \tag{A.21}$$

where κ_t is a stochastic process adapted to $\mathcal{F}$. The optimal wealth and consumption processes take form

$$W_t^\star = G_t \varpi_{1,t}, \quad C_t^\star = G_t \varpi_{2,t}, \tag{A.22}$$

where $\varpi_{1,t}$ and $\varpi_{w,t}$ are adapted to $\mathcal{F}$, and therefore

$$J_t^\star = G_t^{1-\gamma} \varpi_{1,t}^{1-\gamma} \kappa_t. \tag{A.23}$$

We next use these expressions to evaluate the partial derivatives of the aggregator h:

$$\frac{\partial\phi(C_t^\star, J_t^\star; \bar{C}_t^\star)}{\partial J_t^\star} = -\frac{\rho(1-\gamma)}{1-\theta^{-1}} - \frac{\rho(1-\gamma)^{\frac{1-\theta^{-1}}{\gamma-1}}(\gamma-\theta^{-1})}{1-\theta^{-1}} \left(\varpi_{2,t}^{1-h}\left(\varpi_{2,t}/\bar{C}_t\right)^h\right)^{1-\theta^{-1}} \left(\varpi_{1,t}^{1-\gamma}\kappa_t\right)^{\frac{1-\theta^{-1}}{\gamma-1}}, \tag{A.24}$$

which is adapted to $\mathcal{F}$; and

$$\frac{\partial\phi(C_t^\star, J_t^\star; \bar{C}_t)}{\partial C_t^\star} = \rho(1-\gamma)^{\frac{1-\theta^{-1}}{\gamma-1}} \bar{C}_t^{h(1-\theta^{-1})} \varpi_{2,t}^{-\theta^{-1}} \left(\varpi_{1,t}^{1-\gamma}\kappa_t\right)^{\frac{1-\theta^{-1}}{\gamma-1}} G_t^{-\gamma}. \tag{A.25}$$

Thus, the process

$$A_t = \frac{\partial\phi(C_t^\star, J_t^\star; \bar{C}_t)}{\partial C_t^\star} e^{-\gamma\delta t} G_t^\gamma \tag{A.26}$$

is adapted to $\mathcal{F}$. Based on the above results, we express π'_t as

$$\pi'_t = \exp\left(\int_0^t \frac{\partial\phi(C_s, J_s; \bar{C}_s)}{\partial J_s} + \gamma\delta\, ds\right) A_t G_t^{1-\gamma} \tag{A.27}$$

Define $\pi_t = \mathrm{E}[\pi'_t|\mathcal{F}_t]$. Since all asset price processes in the original market are adapted to $\mathcal{F}$, π_t is also a valid state-price density process. Since

$$\mathrm{E}\left[G_t^{1-\gamma}|\mathcal{F}_t\right] = \exp\left(\int_0^t \delta(1-\gamma) + \lambda\left(e^{(1-\gamma)v_s} - 1\right) ds\right), \tag{A.28}$$

we find

$$\pi_t = \exp\left(\int_0^t \frac{\partial\phi(C^\star{}_s, J^\star{}_s; \bar{C}_s)}{\partial J^\star{}_s} + \delta + \lambda\left(e^{(1-\gamma)v_s} - 1\right) ds\right) A_t. \tag{A.29}$$

The last part of the proof consists of proving the fact that

$$\mathrm{E}\left[(G_t)^a|\mathcal{F}_t\right] = \exp\left(\int_0^t \lambda\left(e^{av_s} - 1\right) ds\right) \tag{A.30}$$

To do so, fix the path of v_s, and worry only about uncertainty associated with jumps. Define

$$M_t = \exp\left(\int_0^t v_s a\, dN_s - \int_0^t \lambda\left(e^{av_s} - 1\right)\, ds\right) \tag{A.31}$$

Then

$$dM_t = -M_t \lambda\left(e^{av_s} - 1\right)\, dt + \left(e^{av_s} - 1\right) M_t\, dN_t, \tag{A.32}$$

and therefore

$$\mathrm{E}\left[dM_t | N_s,\, s \leq t\right] = 0 \tag{A.33}$$

and M_t is a martingale. So, $\mathrm{E}\left[M_t|\mathcal{F}_t\right] = 1$. ■

Lemma 3 (Aggregate state) *The vector $X_t = (x_t, \xi_t, K_t)$ is a sufficient statistic for the aggregate state of the economy.*

Proof. We first establish that X_t is a Markov process. The optimal investment in a new project j at time t is the solution to

$$\nu_t \equiv \max_{k_{j,t}} \mathrm{E}_t\left[\int_t^\infty \frac{\pi_s}{\pi_t}\, p_{Z,s}\, \zeta_{j,s}\, ds\right] - k_{j,t}^{1/\alpha} = P_t\, k_{j,t} e^{\xi_t} - k_{j,t}^{1/\alpha}, \tag{A.34}$$

where

$$P_t = \mathrm{E}_t\left[\int_t^\infty \frac{\pi_s}{\pi_t} e^{-\delta(s-t)}\, p_{Z,t}\, ds\right]. \tag{A.35}$$

The optimal scale of each new project is

$$k_t^\star = \left(\alpha\, e^{\xi_t}\, P_t\right)^{\frac{\alpha}{1-\alpha}}. \tag{A.36}$$

Clearing the market for new capital yields

$$k_t^\star = \left(\frac{I_t}{\lambda}\right)^\alpha. \tag{A.37}$$

As a result, the NPV of a new project equals

$$\nu_t = \left(\alpha^{\alpha/(1-\alpha)} - \alpha\right)\left(P_t e^{\xi_t}\right)^{1/(1-\alpha)}. \tag{A.38}$$

The fact that investment decisions are independent of u also implies that $Z_t = K_t$. The evolution of K is given by

$$dK_t = \left(-\delta\, K_t + \lambda\, e^{\xi_t}\left(\frac{I_t}{\lambda}\right)^\alpha\right) dt. \tag{A.39}$$

Next, we need to show that prices $(\mathrm{w}_t, p_{Z,t}, \pi_t)$ depend only on the aggregate state X. To see that this holds for the first two, consider the profit maximizing conditions for the final good firm,

$$p_{Zt} = \phi\, e^{(1-\phi)\, x_t}\, K_t^{\phi-1} L_t^{1-\phi} \tag{A.40}$$

$$\mathrm{w}_t = (1-\phi)\, e^{(1-\phi)\, x_t}\, K_t^{\phi} L_t^{-\phi}. \tag{A.41}$$

What remains is to show that π_t is independent of the cross-sectional distribution of household wealth. The fact that preferences are homothetic implies that it suffices to show that the growth rate of a household's wealth is independent of wealth. The evolution of a household's wealth share is given by

$$\frac{dw_{n,t}}{w_{n,t}} = \delta^h\, dt - \frac{\lambda\eta\nu_t}{V_t + G_t + \psi\, H_t}\, dt + \psi\, \frac{\lambda\eta\nu_t}{V_t + G_t + \psi\, H_t}\, \mu_I^{-1} dN_{it}^I. \tag{A.42}$$

The benchmark model in the paper corresponds to the case where $\psi = 1$. Here,

$$V_t \equiv \int_{j\in\mathcal{J}_t} \left(\mathrm{E}_t \left[\int_t^\infty \frac{\pi_s}{\pi_t} \left(p_{Z,s}\, \zeta_{j,s}\, ds \right) \right] \right) dj = P_t\, K_t, \tag{A.43}$$

$$G_t \equiv (1-\eta) \int_0^1 \left(\mathrm{E}_t \left[\int_t^\infty \frac{\pi_s}{\pi_t} \left(\lambda_{fs} \nu_s \right) ds \right] \right) df = (1-\eta)\, \mathrm{E}_t \left[\int_t^\infty \frac{\pi_s}{\pi_t} \lambda \nu_s\, ds \right], \tag{A.44}$$

$$H_t \equiv \mathrm{E}_t \left[\int_t^\infty e^{-\delta^h (s-t)} \frac{\pi_s}{\pi_t} \mathrm{w}_s\, ds \right]. \tag{A.45}$$

Here, note the role of incomplete markets: households earn an additional non-tradeable return when innovating (upon the arrival of the process dN_{it}^I). However due to the assumption that the measure of new projects is proportional to a household's existing wealth, $w_{n,t}$ does not appear directly in the right hand side of (A.42). The functional form of preferences (12-13) together with our assumption that the scale of the household-level innovation process is proportional to individual wealth imply that optimal individual consumption and portfolio plans are proportional to individual wealth. Then, a household's consumption share is the same as its wealth share, $c_{n,t} = w_{n,t}$. As a result, π_t only depends on the aggregate state vector X_t. ■

Lemma 4 (Investor value function) *The value function of an investor with relative wealth w is given by*

$$J(w, z, \omega) = \frac{1}{1-\gamma} w^{(1-\gamma)} e^{(1-\gamma_1) z}\, f(\omega) \tag{A.46}$$

where z and ω are defined in equations (A.64) and (A.63) in the main text, the function f solves the ODE

$$\begin{aligned} -A_1(\omega) f(\omega)^{\frac{\gamma-\theta^{-1}}{\gamma-1}} =& f(\omega) \left\{ c_0^f - (1-\gamma)\, s(\omega) + \left[\left(1 + \frac{\psi}{\mu_I}\, s(\omega) \right)^{1-\gamma} - 1 \right] \mu_I \right\} \\ &+ f'(\omega) \left\{ c_1^f - (1-\alpha\,\phi)\, \kappa(\omega) \right\} + f''(\omega)\, c_2^f \end{aligned} \tag{A.47}$$

the functions $s(\omega)$ and $\kappa(\omega)$ are given by

$$s(\omega) = \frac{\eta\,(1-\alpha)\, v(\omega)\, \kappa(\omega)}{v(\omega) + g(\omega) + \psi\, h(\omega)} \tag{A.48}$$

$$\kappa(\omega) = \lambda^{1-\alpha}\, e^{(1-\alpha\phi)\,\omega} [i(\omega)]^\alpha, \tag{A.49}$$

where

$$\psi = \frac{\mu_I + q_S\, \delta^h}{\mu_I + \delta^h} \tag{A.50}$$

is the fraction of households participating in the stock market. The functions $v(\omega)$, $g(\omega)$ and $h(\omega)$ solve the system of ODEs

$$\begin{aligned} 0 =& A_v\Big(\omega, i(\omega)\Big) B(\omega) + v'(\omega) \left\{ c_1^f - (1-\alpha\,\phi)\, \kappa(\omega) \right\} + v''(\omega)\, c_2^f \\ &+ v(\omega) \left\{ c_0^f - \frac{\gamma - \theta^{-1}}{1-\gamma} A_1(\omega) f(\omega)^{\frac{1-\theta^{-1}}{\gamma-1}} + \mu_I \left(\left(1 + \frac{\psi}{\mu_I} s(\omega) \right)^{1-\gamma} - 1 \right) + \gamma\, s(\omega) - \kappa(\omega) \right\} \\ 0 =& (1-\eta)\,(1-\alpha)\, v(\omega)\, \kappa(\omega) + g'(\omega) \left\{ c_1^f - (1-\alpha\,\phi)\, \kappa(\omega) \right\} + g''(\omega)\, c_2^f \end{aligned} \tag{A.51}$$

$$+ g(\omega)\left\{ c_0^f - \frac{\gamma-\theta^{-1}}{1-\gamma} A_1(\omega) f(\omega)^{\frac{1-\theta^{-1}}{\gamma-1}} + \mu_I \left(\left(1+\frac{\psi}{\mu_I} s(\omega)\right)^{1-\gamma} - 1 \right) + \gamma\, s(\omega) \right\} \tag{A.52}$$

$$0 = A_h\Big(\omega, i(\omega)\Big) B(\omega) + h'(\omega)\left\{ c_1^f - (1-\alpha\,\phi)\,\kappa(\omega) \right\} + h''(\omega)\, c_2^f$$

$$+ h(\omega)\left\{ c_0^f - \delta^h - \frac{\gamma-\theta^{-1}}{1-\gamma} A_1(\omega) f(\omega_t)^{\frac{1-\theta^{-1}}{\gamma-1}} + \mu_I \left(\left(1+\frac{\psi}{\mu_I} s(\omega_t)\right)^{1-\gamma} - 1 \right) + \gamma\, s(\omega) \right\} \tag{A.53}$$

where the constants c_0^f, c_1^f, c_2^f *and* ϕ^d *are*

$$c_0^f = \Bigg\{ \delta^h(1-\gamma) - \frac{\rho(1-\gamma)}{1-\theta^{-1}} + (1-\gamma_1)(1-\phi)\mu_x + \frac{1}{2}(1-\phi)^2\,\sigma_x^2\,(1-\gamma_1)^2$$

$$+ \frac{\phi(1-\gamma_1)}{1-\alpha\,\phi}\left(\mu_\xi + \alpha\,(1-\phi)\,\mu_x + (1-\gamma_1)\,\alpha(1-\phi)^2\,\sigma_x^2 \right) + \frac{1}{2}\left(\frac{\phi(1-\gamma_1)}{1-\alpha\,\phi}\right)^2 \left(\sigma_\xi^2 + \alpha^2(1-\phi)^2\sigma_x^2\right) \Bigg\} \tag{A.54}$$

$$c_1^f = \left\{ \mu_\xi + \alpha\,(1-\phi)\,\mu_x + (1-\alpha\,\phi)\delta + (1-\gamma_1)\,\alpha(1-\phi)^2\,\sigma_x^2 + \frac{\phi(1-\gamma_1)}{1-\alpha\,\phi}\,\left(\sigma_\xi^2 + \alpha^2(1-\phi)^2\sigma_x^2\right) \right\} \tag{A.55}$$

$$c_2^f = \frac{1}{2}\,\left(\sigma_\xi^2 + \alpha^2(1-\phi)^2\sigma_x^2\right) \tag{A.56}$$

$$\tag{A.57}$$

and the following functions are defined as

$$A_1(\omega) = \frac{\rho(1-\gamma)}{1-\theta^{-1}}\,\left[(1-i(\omega))\right]^{1-\theta_1^{-1}} \left(1 - \frac{(1-\psi)(1-\phi)}{1-i(\omega)}\right)^{1-\theta^{-1}} e^{-\phi(1-\theta_1^{-1})\,\omega} \tag{A.58}$$

$$A_v(\omega) = \phi\, e^{-\phi\omega} \tag{A.59}$$

$$A_h(\omega) = (1-\phi)\, e^{-\phi\omega} \tag{A.60}$$

$$B(\omega) = \rho\,\left[(1-i(\omega))\right]^{-\theta_1^{-1}} \left(1 - \frac{(1-\psi)(1-\phi)}{1-i(\omega)}\right)^{-1/\theta} f(\omega)^{\frac{\gamma-\theta^{-1}}{\gamma-1}}\, e^{\phi\,(1-\theta_1^{-1})\,\omega} \tag{A.61}$$

and the function $i(\omega)$ *is the solution to*

$$\alpha\, v(\omega) f(\omega)^{\frac{\gamma-\theta^{-1}}{1-\gamma}}\, \exp\Big(\left(1-\alpha\phi-\phi\,(1-\theta_1^{-1})\right)\omega\Big)\,\left[(1-i(\omega))\right]^{\theta_1^{-1}} \left(1 - \frac{(1-\psi)(1-\phi)}{1-i(\omega)}\right)^{1/\theta} = \left(\frac{i(\omega)}{\lambda}\right)^{1-\alpha}. \tag{A.62}$$

Proof. We first show that the

$$\omega_t = \Big(\xi_t + \alpha\,\chi_t - \log K_t\Big). \tag{A.63}$$

and

$$\chi_t = \frac{1-\phi}{1-\alpha\,\phi}\,x_t + \frac{\phi}{1-\alpha\,\phi}\xi_t. \tag{A.64}$$

are sufficient to characterize the state vector X_t. We first conjecture, and subsequently verify, that

$$I_t = i(\omega_t)\, e^{(1-\phi)\,x_t}\, K_t^{\phi}. \tag{A.65}$$

Using equations (A.36)-(A.37) above, we see that the function $i(\omega_t)$ solves:

$$\left(i(\omega_t)\, e^{(1-\phi)\, x_t}\, K_t^{\phi}\right)^{1-\alpha} = \alpha e^{\xi_t} P(X_t) \lambda^{1-\alpha} \tag{A.66}$$

Using our conjecture (A.65) and market clearing (A.37), we see that the evolution of K_t depends only on ω,

$$\frac{dK_t}{K_t} = -\delta\, dt + \lambda \frac{e^{\xi_t}\left(\lambda^{-1} e^{(1-\phi)\, x_t}\, K_t^{\phi}\, i(\omega_t)\right)^{\alpha}}{K_t} = -\delta\, dt + \kappa(\omega_t)\, dt, \qquad \kappa(\omega_t) \equiv \lambda\, e^{(1-\alpha\,\phi)\,\omega_t} \left(\frac{i(\omega_t)}{\lambda}\right)^{\alpha}$$

Next, we consider the dynamics of the SDF π. Consider the evolution of household's wealth share (A.42). We first conjecture – and subsequently verify – that the following ratio is only a function of ω

$$s(\omega_t) \equiv \frac{\lambda \eta \nu_t}{V_t + G_t + \psi\, H_t}. \tag{A.67}$$

Using Lemma 2 above, we get that

$$\pi_t = \exp\left(\int_0^t \delta^h + \frac{\partial \phi(C_s^{\star}, J_s^{\star}; \bar{C}_s)}{\partial J_s^{\star}} + \mu_I \left(\left(1 + \frac{\psi}{\mu_I} s(\omega_s)\right)^{1-\gamma} - 1\right) ds\right) A_t, \tag{A.68}$$

where

$$A_t \equiv \frac{\partial \phi(C_t^{\star}, J_t^{\star}; \bar{C}_t)}{\partial C_t^{\star}} \exp\left(\int_0^t \gamma \log\left(1 + \frac{\psi}{\mu_I} s(\omega_s)\right) dN_s\right) \tag{A.69}$$

and we have applied Lemma 2 using

$$v_t \equiv \log\left(1 + \frac{\psi}{\mu_I} s(\omega_t)\right). \tag{A.70}$$

Using equations (A.40)-(A.41), along with the definition of ω and χ implies that the consumption of shareholders as a group is

$$\begin{aligned} C_t^S &= C_t - (1-\psi)\, \mathrm{w}_t \\ &= e^{\chi_t} e^{-\phi\, \omega_t} \left(1 - i(\omega_t) - (1-\psi)(1-\phi)\right). \end{aligned} \tag{A.71}$$

Recall that due to homotheticity implies that the consumption of a household with wealth share j is

$$C_{j,t} = w_{j,t}\, C_t^S. \tag{A.72}$$

To find the SDF, we first need to find the value function J of a household. Conjecture that the value function is given by (A.46).This guess needs to satisfy the Hamilton-Jacobi-Bellman (HJB) equation,

$$\phi(C, J; \bar{C}) + \mathcal{D}_{X,w} J = 0. \tag{A.73}$$

Plugging the guess into the HJB equation, we get an ODE in $f(\omega)$ – equation (A.47) in the proposition. Thus, if a solution to equation (A.47) exists (and is bounded) equation (A.46) is the household's value function.

What remains is to show that our conjectures for $i(\omega)$ and $s(\omega)$ are correct. Using (A.46) and the evolution of shareholder consumption (A.72), we get that the last term in the SDF – given by (A.69) – equals

$$A_t = e^{-\gamma_1\, \chi_t} \hat{w}_t^{-\gamma}\, B(\omega_t), \tag{A.74}$$

where $B(\omega)$ is defined in (A.61) and the process $\hat{w}_t$ is related to the relative wealth share, but without the idiosyncratic jumps

$$d\hat{w}_t = \left(\delta^h - s(\omega_t)\right)\, dt. \tag{A.75}$$

Here, we note that all the terms inside the integral in (A.68) are just a function of ω,

$$\frac{\partial \phi(C_t^\star, J_t^\star; \bar{C}_t)}{\partial J_t^\star} = -\frac{\rho}{1-\theta^{-1}} \left((\gamma - \theta^{-1})\, l(\omega)^{1-\theta^{-1}} \, [f(\omega)]^{\frac{1-\theta^{-1}}{\gamma-1}} + (1-\gamma) \right) \tag{A.76}$$

where

$$l(\omega) \equiv (1 - i(\omega) - (1-\psi)(1-\phi)) \, ((1 - i(\omega)))^{-h} \tag{A.77}$$

Next, we compute the value of assets in place. Using the SDF (A.68), equation (A.40), the definition of z and ω and the Feynman-Kac (FK) theorem with stochastic discounting, we have that

$$P(X_t) = K_t^{-1} e^{\chi_t} \, v(\omega_t) B(\omega_t)^{-1} \tag{A.78}$$

where $v(\omega)$ solves the ODE (A.51). Using (A.78) with (A.66) yields equation (A.62) in the proposition. It also confirms our conjecture above that $i(\omega)$ is only a function of ω. We next verify that the same is true for $s(\omega_t)$. The value of assets in place in the economy is equal to

$$V_t = e^{\chi_t} \, v(\omega_t) B(\omega_t)^{-1}. \tag{A.79}$$

Next, we solve for the value of growth opportunities. Using the optimal investment policy, we get that the net present value of a new project is equal to

$$\nu_t = (1-\alpha) \, e^{\chi_t} \, v(\omega_t) \, B(\omega_t)^{-1} \, e^{(1-\alpha\,\phi)\,\omega_t} \left(\frac{i(\omega_t)}{\lambda} \right)^{\alpha} \tag{A.80}$$

The value of growth opportunities for the average firm ($\lambda_f = \lambda$) equals

$$G_t = e^{\chi_t} g(\omega_t) \, (B(\omega_t))^{-1} \tag{A.81}$$

where $g(\omega)$ solves the ODE in (A.52). As before, we have used the FK theorem, along with the definition of z and ω. The last step consists of computing the value of human capital. Using the equilibrium wage (A.41), the FK theorem, along with the definition of z and ω, we get that

$$H_t = e^{\chi_t} h(\omega_t) \, (B(\omega_t))^{-1}, \tag{A.82}$$

where $h(\omega)$ solves the ODE in (A.53). Equations (A.79)-(A.80), (A.81)-(A.82) imply that $s(\omega_t)$ is only a function of ω and is equal to equation (A.48) in the proposition. This completes the proof.

■

Proof of Proposition 1. Proposition follows directly from (A.68) and (A.74) in the proof of the lemma above. The process $b(\omega_t)$ is given by

$$b(\omega) = (1-\gamma)\delta^h - \rho\,\kappa - \rho\,(1-\kappa) \, (1 - i(\omega_t))^{(1-\theta_1^{-1})} \, (f(\omega))^{-\kappa^{-1}} + \gamma\, s(\omega) + \mu_I \left(\left(1 + \frac{\psi}{\mu_I} s(\omega) \right)^{1-\gamma} - 1 \right). \tag{A.83}$$

where the functions $i(\omega)$, $f(\omega)$, and $s(\omega)$ are defined in the proof of lemma above. ■

Proposition 2 (Market value of a firm) *The market value of a firm equals*

$$S_{f,t} = e^{\chi_t} \left[v(\omega_t) \, \frac{z_{f,t}}{\bar{u}_{ft}} \left(1 + \frac{v_1(\omega_t)}{v(\omega_t)} \, (\bar{u}_{f,t} - 1) \right) + g(\omega_t) + \left(\frac{\lambda_{f,t}}{\lambda} - 1 \right) g_1(\omega_t) \right] (B(\omega_t))^{-1} \tag{A.84}$$

where $v(\omega)$ and $g(\omega)$ are defined above and the functions v_1 and g_1 solve the ODEs

$$0 = A_v(\omega)B(\omega) + v_1'(\omega)\left\{c_1^f - (1-\alpha\,\phi)\,\kappa(\omega)\right\} + v_1''(\omega)\,c_2^f$$
$$+\, v_1(\omega)\left\{c_0^f - \kappa_u - \frac{\gamma-\theta^{-1}}{1-\gamma}A_1(\omega)f(\omega)^{\frac{1-\theta^{-1}}{\gamma-1}} + \mu_I\left(\left(1+\frac{\psi}{\mu_I}s(\omega)\right)^{1-\gamma} - 1\right) + \gamma\,s(\omega) - \kappa(\omega)\right\} \quad \text{(A.85)}$$

$$0 = (1-\eta)\,(1-\alpha)\,v(\omega)\,\kappa(\omega) + g_1'(\omega)\left\{c_1^f - (1-\alpha\,\phi)\,\kappa(\omega)\right\} + g_1''(\omega)\,c_2^f$$
$$+\, g_1(\omega)\left\{c_0^f - \mu_L - \mu_H - \frac{\gamma-\theta^{-1}}{1-\gamma}A_1(\omega)f(\omega)^{\frac{1-\theta^{-1}}{\gamma-1}} + \mu_I\left(\left(1+\frac{\psi}{\mu_I}s(\omega)\right)^{1-\gamma} - 1\right) + \gamma\,s(\omega)\right\} \quad \text{(A.86)}$$

Proof. The proof follows closely the derivations of equations (A.78) and (A.81) above. We have that the value of assets in place for an existing firm with capital stock $K_{f,t}$ and profitability $Z_{f,t}$ are given by

$$VAP_{f,t} = P(X_t)K_{f,t} + P_1(X_t)\,(Z_{f,t} - K_{f,t})\,, \quad \text{(A.87)}$$

where

$$P_1(X_t) \equiv \mathrm{E}_t\left[\int_t^\infty \frac{\pi_s}{\pi_t}e^{-(\delta+\kappa_u)(s-t)}\,p_{Z,s}\,ds\right] = K_t^{-1}e^{\chi_t}\,v_1(\omega_t)B(\omega_t)^{-1} \quad \text{(A.88)}$$

where $v_1(\omega)$ satisfies the ODE (A.85). As above, we have used the SDF (A.68), equation (A.40), the definition of z and ω and the Feynman-Kac (FK) theorem with stochastic discounting. Similarly, the present value of growth opportunities for a firm equals

$$PVGO_{f,t} \equiv (1-\eta)\,E_t\int_t^\infty \frac{\pi_s}{\pi_t}\lambda_{f,s}\,\nu_s\,ds$$
$$= PVGO_t + \lambda\,(1-\eta)\left(\frac{\lambda_{f,t}}{\lambda} - 1\right)E_t\int_t^\infty \frac{\pi_s}{\pi_t}e^{(\mu_L+\mu_H)(s-t)}\,\nu_s\,ds \quad \text{(A.89)}$$
$$= e^{\chi_t}g(\omega_t)\,(B(\omega_t))^{-1} + e^{\chi_t}g_1(\omega_t)\,(B(\omega_t))^{-1} \quad \text{(A.90)}$$

where $g_1(\omega)$ satisfies the ODE (A.85). As above, we have used the SDF (A.68), the definition of z and ω and the Feynman-Kac (FK) theorem with stochastic discounting, and the fact that

$$E[\lambda_{fs}|\lambda_{f,t}] = \lambda + \lambda\left(\frac{\lambda_{f,t}}{\lambda} - 1\right)e^{(\mu_L+\mu_H)(s-t)}. \quad \text{(A.91)}$$

■

Appendix B: Data

Construction of Estimation Targets

Aggregate consumption: We report the mean and volatility of log consumption growth. We use the Barro and Ursua (2008) data for the United States. We compute the estimate of LR variability using the estimator in Dew-Becker (2014).

Volatility of shareholder consumption growth: The volatility of shareholder consumption growth is from the unpublished working paper version of Malloy et al. (2009) and includes their

adjustment for measurement error. We are grateful to Annette Vissing-Jorgensen for suggesting this.

Aggregate Investment and Output: Investment is non-residential private domestic investment. Output is gross domestic product. Both series are deflated by population and the CPI. Data on the CPI are from the BLS. Population is from the Census Bureau.

Dividends and Payout: Moments of net payout to assets are from Larrain and Yogo (2008).

Firm Investment rate and Q: Firm investment is defined as the change in log gross PPE. Tobin's Q equals the market value of equity (CRSP December market cap) plus book value of preferred shares plus long term debt minus inventories and deferred taxes over book assets. In the model, we construct Tobin's Q as the ratio of the market value of the firm divided by the replacement cost of capital. Replacement cost is defined as the current capital stock valued at the price of new capital, adjusted for quality. When computing correlation coefficients, we winsorize the data by year at the 1% level to minimize the effect of outliers.

Firm profitability: Firm profitability equals gross profitability (sales minus costs of goods sold) scaled by book capital (PPE). In the model, we compute firm profitability as $p_{Y,t}Y_{f,t}/(p_{I,t}K_{f,t})$. When computing correlation coefficients, we winsorize the data by year at the 1% level to minimize the effect of outliers.

Market portfolio and risk-free rate moments: Using reported estimate from the long sample of Barro and Ursua (2008) for the United States.

Value factor moments: We use the data from Kenneth French. Value factor is the 10m1 portfolio of firms sorted on book-to-market.

Consumption share of stockholders: Consumption share of stock holders is from Table 2 of Giuvenen 2006. This number is also consistent with estimates using the PSID. This number is also consistent with Heaton and Lucas (2000): using their data on Table AII we get an income share for stockholders of about 43%.

Relative consumption growth of shareholders: We use the series constructed in Malloy et al. (2009), which covers the 1982-2004 period. We follow Jagannathan and Wang (2007) and construct annual consumption growth rates by using end-of-period consumption. In particular, we focus on the sample of households that are interviewed in December of every year, and use the average 8 quarter consumption growth rate of non-stockholders and stockholders, defined as in Malloy et al. (2009). We focus on 2-yr horizon. Results using longer horizons are available upon request.

Details on the construction of the innovation measure

We create $\hat{\omega}$ using a non-parametric variant of the Kogan et al. (2012) procedure. First, we create idiosyncratic stock returns for firm f around the day that patent j is granted to equal the 3-day return of the firm minus the return on the CRSP value-weighted index around the same window,

$$r^e_{fj} = r_{fj} - r_{mj}. \tag{B.1}$$

Patents are issued every Tuesday. Hence, r_{fj} are the accumulated return over Tuesday, Wednesday and Thursday following the patent issue. Second, we compute an estimate of the value of patent j as the firm's market capitalization on the day prior the patent announcement V_{fj} times the idiosyncratic return to the firm truncated at zero,

$$\hat{\nu}_j = \frac{1}{N_j} \max(r^e_{fj}, 0)\, V_{fj}. \tag{B.2}$$

If multiple patents were granted in the same day to the same firm, we divide by the number of patents N. Relative to Kogan et al. (2012), we replace the filtered value of the patent $E[x_{fd}|r_{fd}]$ with $\max(r^e_{fj}, 0)$. Our construction is an approximation to the measure in Kogan et al. (2012) that can be easily implemented in simulated data without the additional estimation of parameters. Third, we construct an estimate of NPV created at the firm level as

$$\hat{\nu}_{ft} = \sum_{j \in J_{ft}} \hat{\nu}_j, \tag{B.3}$$

where J_{ft} is the set of patents granted to firm f in year t. Last, to construct an estimate of the aggregate state ω, we sum across the estimated patent values at year t, and scale by total market capitalization

$$\hat{\omega}_t \equiv \frac{\sum_{f \in F_t} \hat{\nu}_{ft}}{\sum_{f \in F_t} V_{ft}}, \tag{B.4}$$

where F_t is the set of firms in year t and V_{ft} is the market capitalization of firm f at the end of year t. We compute $\hat{\omega}_{I \setminus f}$ as in (B.4), except that the summation is over all firms in the same industry (defined at the 3-digit SIC code level) excluding firm f.

We follow a similar approach when constructing $\hat{\omega}$ in simulated data. We compute the excess return of the firm as in equation (B.1) around the stochastic times that the firm acquires a new project, and then construct $\hat{\nu}_j$ as in equation (B.2). We then compute

References

Abel, A. B. (1983). Optimal investment under uncertainty. *American Economic Review 73*(1), 228–33.

Abel, A. B. (1990). Asset prices under habit formation and catching up with the joneses. *American Economic Review 80*(2), 38–42.

Abel, A. B. (1999, February). Risk premia and term premia in general equilibrium. *Journal of Monetary Economics 43*(1), 3–33.

Ai, H., M. M. Croce, and K. Li (2013). Toward a quantitative general equilibrium asset pricing model with intangible capital. *Review of Financial Studies 26*, 491–530.

Autor, D. H., L. F. Katz, and M. S. Kearney (2006, May). The Polarization of the U.S. Labor Market. *American Economic Review 96*(2), 189–194.

Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance 59*(4), pp. 1481–1509.

Barro, R. J. and J. F. Ursua (2008). Macroeconomic Crises since 1870. *Brookings Papers on Economic Activity 39*(1 (Spring), 255–350.

Basu, S., J. G. Fernald, and M. S. Kimball (2006). Are technology improvements contractionary? *American Economic Review 96*(5), 1418–1448.

Berk, J. B., R. C. Green, and V. Naik (1999). Optimal investment, growth options, and security returns. *Journal of Finance 54*(5), 1553–1607.

Berndt, E. (1990). Energy use, technical progress and productivity growth: A survey of economic issues. *Journal of Productivity Analysis 2*(1), 67–83.

Björkman, M. and K. Holmström (2000). Global optimization of costly nonconvex functions using radial basis functions. *Optimization and Engineering 1*(4), 373–397.

Blanchard, O. J. (1985). Debt, deficits, and finite horizons. *Journal of Political Economy 93*(2), 223–47.

Boldrin, M., L. J. Christiano, and J. D. M. Fisher (2001). Habit persistence, asset returns, and the business cycle. *American Economic Review 91*(1), 149–166.

Bolton, P., N. Wang, and J. Yang (2015). A theory of liquidity and risk management based on the inalienability of risky human capital. Working paper, Columbia University.

Campanale, C., R. Castro, and G. L. Clementi (2010). Asset pricing in a production economy with chew-dekel preferences. *Review of Economic Dynamics 13*(2), 379–402.

Carlson, M., A. Fisher, and R. Giammarino (2004). Corporate investment and asset price dynamics: Implications for the cross-section of returns. *The Journal of Finance 59*(6), pp. 2577–2603.

Clark, A. E. and A. J. Oswald (1996). Satisfaction and comparison income. *Journal of Public Economics 61*(3), 359 – 381.

Constantinides, G. M. and D. Duffie (1996, April). Asset Pricing with Heterogeneous Consumers. *Journal of Political Economy 104*(2), 219–40.

Constantinides, G. M. and A. Ghosh (2014, May). Asset Pricing with Countercyclical Household Consumption Risk. NBER Working Papers 20110, National Bureau of Economic Research, Inc.

Cooley, T. F., J. Greenwood, and M. Yorukoglu (1997). The replacement problem. *Journal of Monetary Economics 40*(3), 457–499.

Croce, M. M. (2014). Long-run productivity risk: A new hope for production-based asset pricing? *Journal of Monetary Economics 66*(0), 13 – 31.

DeMarzo, P. M., R. Kaniel, and I. Kremer (2008). Relative wealth concerns and financial bubbles. *Review of Financial Studies 21*(1), 19–50.

Dew-Becker, I. (2014). How risky is consumption in the long-run? benchmark estimates from a novel unbiased and efficient estimator. Working paper, Northwestern University.

Duesenberry, J. (1949). *Income, saving, and the theory of consumer behavior*. Harvard economic studies. Harvard University Press.

Duffie, D. and L. G. Epstein (1992). Stochastic differential utility. *Econometrica 60*(2), 353–94.

Epstein, L. G. and S. E. Zin (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica 57*(4), 937–69.

Fama, E. F. and K. R. French (1992). The cross-section of expected stock returns. *Journal of Finance 47*(2), 427–65.

Fama, E. F. and K. R. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics 33*(1), 3–56.

Fama, E. F. and K. R. French (1995). Size and book-to-market factors in earnings and returns. *Journal of Finance 50*(1), 131–55.

Ferrer-i Carbonell, A. (2005). Income and well-being: an empirical analysis of the comparison income effect. *Journal of Public Economics 89*(5-6), 997–1019.

Fisher, J. D. M. (2006). The dynamic effects of neutral and investment-specific technology shocks. *Journal of Political Economy 114*(3), 413–451.

Frydman, C. (2015). What drives peer effects in financial decision-making? neural and behavioral evidence. Working paper, USC.

Gali, J. (1994). Keeping up with the joneses: Consumption externalities, portfolio choice, and asset prices. *Journal of Money, Credit and Banking 26*(1), 1–8.

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

Garleanu, N. and S. Panageas (2014). Young,old,conservative and bold:the implications of heterogeneity and finite lives for asset pricing. Technical report.

Garleanu, N., S. Panageas, and J. Yu (2012). Technological growth and asset pricing. *Journal of Finance 67*, 1265–1292.

Gomes, J. F., L. Kogan, and L. Zhang (2003). Equilibrium cross section of returns. *Journal of Political Economy 111*(4), 693–732.

Greenwood, J., Z. Hercowitz, and P. Krusell (1997). Long-run implications of investment-specific technological change. *American Economic Review 87*(3), 342–362.

Guvenen, F. (2006). Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective. *Journal of Monetary Economics 53*(7), 1451 – 1472.

Guvenen, F. (2009, November). A Parsimonious Macroeconomic Model for Asset Pricing. *Econometrica 77*(6), 1711–1750.

Hansen, L. P., J. C. Heaton, and N. Li (2005). Intangible risk. In *Measuring Capital in the New Economy*, NBER Chapters, pp. 111–152. National Bureau of Economic Research, Inc.

Hart, O. and J. Moore (1994, November). A Theory of Debt Based on the Inalienability of Human Capital. *The Quarterly Journal of Economics 109*(4), 841–79.

Hurst, E. and A. Lusardi (2004). Liquidity constraints, household wealth, and entrepreneurship. *Journal of Political Economy 112*(2), pp. 319–347.

Ingram, B. F. and B.-S. Lee (1991, February). Simulation estimation of time-series models. *Journal of Econometrics 47*(2-3), 197–205.

Jagannathan, R. and Y. Wang (2007). Lazy investors, discretionary consumption, and the cross-section of stock returns. *Journal of Finance 62*(4), 1623–1661.

Jensen, J. B., R. H. McGuckin, and K. J. Stiroh (2001). The impact of vintage and survival on productivity: Evidence from cohorts of u.s. manufacturing plants. *The Review of Economics and Statistics 83*(2), 323–332.

Jermann, U. J. (1998). Asset pricing in production economies. *Journal of Monetary Economics 41*(2), 257–275.

Kaltenbrunner, G. and L. A. Lochstoer (2010). Long-run risk through consumption smoothing. *Review of Financial Studies 23*(8), 3190–3224.

Kogan, L. and D. Papanikolaou (2012). Economic activity of firms and asset prices. *Annual Review of Financial Economics 4*(1).

Kogan, L. and D. Papanikolaou (2013). Firm characteristics and stock returns: The role of investment-specific shocks. *Review of Financial Studies 26*(11).

Kogan, L. and D. Papanikolaou (2014). Growth opportunities, technology shocks, and asset prices. *Journal of Finance 69*(2).

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2012). Technological innovation, resource allocation, and growth. Working paper 17769, National Bureau of Economic Research.

Larrain, B. and M. Yogo (2008). Does firm value move too much to be justified by subsequent changes in cash flow? *Journal of Financial Economics 87*(1), 200 – 226.

Luttmer, E. F. P. (2005). Neighbors as negatives: Relative earnings and well-being. *The Quarterly Journal of Economics 120*(3), 963–1002.

Malloy, C. J., T. J. Moskowitz, and A. Vissing-Jorgensen (2009). Long-run stockholder consumption risk and asset returns. *Journal of Finance 64*(6), 2427–2479.

Papanikolaou, D. (2011). Investment shocks and asset prices. *Journal of Political Economy 119*(4), pp. 639–685.

Pastor, L. and P. Veronesi (2009). Technological revolutions and stock prices. *American Economic Review 99*(4), 1451–83.

Piketty, T. and E. Saez (2003). Income inequality in the united states, 19131998. *The Quarterly Journal of Economics 118*(1), 1–41.

Poterba, J. M. and A. A. Samwick (1995). Stock Ownership Patterns, Stock Market Fluctuations, and Consumption. *Brookings Papers on Economic Activity 26*(2), 295–372.

Rauh, J. D. and A. Sufi (2011). Explaining corporate capital structure: Product markets, leases, and asset similarity. *Review of Finance 16*(1), 115–155.

Rayo, L. and G. S. Becker (2007). Evolutionary Efficiency and Happiness. *Journal of Political Economy 115*, 302–337.

Roussanov, N. (2010, October). Diversification and its discontents: Idiosyncratic and entrepreneurial risk in the quest for social status. *Journal of Finance 65*(5), 1755–1788.

Rouwenhorst, G. K. (1995). *Frontiers of business cycle research*, Chapter Asset pricing implications of equilibrium business cycle models. Princeton University Press.

Solnick, S. and D. Hemenway (1998). Is more always better?: A survey on positional concerns. *Journal of Economic Behavior & Organization 37*(3), 373–383.

Solow, R. M. (1960). Investment and technical progess. In K. J. Arrow, A. Karlin, and P. Suppes (Eds.), *Mathematical Methods in the Social Sciences*, pp. 89–104. Stanford, CA: Stanford University Press.

Storesletten, K., C. Telmer, and A. Yaron (2007, October). Asset Pricing with Idiosyncratic Risk and Overlapping Generations. *Review of Economic Dynamics 10*(4), 519–548.

Zhang, L. (2005). The value premium. *Journal of Finance 60*(1), 67–103.

Table 1: Benchmark Estimates: Goodness of Fit

Statistic	Data	Model			
		Mean	RDev2.	(5%)	(95%)
Aggregate quantities					
Consumption growth, mean	0.015	0.014	0.001	0.004	0.024
Consumption growth, volatility	0.036	0.042	0.028	0.038	0.047
Consumption growth, long-run volatility	0.041	0.056	0.128	0.040	0.072
Investment, mean share of output	0.089	0.055	0.148	0.028	0.083
Investment, volatility	0.130	0.116	0.012	0.094	0.137
Investment and consumption, correlation	0.472	0.383	0.036	0.167	0.554
Net payout to assets, coefficient of variation	0.575	0.520	0.009	0.291	0.850
Asset Prices					
Market Portfolio, excess returns, mean	0.063	0.068	0.005	0.049	0.083
Market Portfolio, excess returns, std	0.185	0.152	0.033	0.138	0.166
Risk-free rate, mean	0.020	0.025	0.065	0.019	0.033
Risk-free rate, volatility	0.007	0.007	0.000	0.003	0.012
Value factor, mean	0.065	0.054	0.031	0.031	0.075
Value factor, volatility	0.243	0.191	0.045	0.134	0.278
Value factor, CAPM alpha	0.040	0.034	0.026	0.007	0.059
Cross-sectional (firm) moments					
Investment rate, IQR	0.175	0.200	0.021	0.153	0.249
Investment rate, serial correlation	0.223	0.191	0.020	0.014	0.429
Investment rate, correlation with Tobin's Q	0.237	0.205	0.019	0.143	0.328
Tobin's Q, IQR	1.139	0.750	0.117	0.534	1.045
Tobin's Q, serial correlation	0.889	0.924	0.002	0.897	0.947
Profitability, IQR	0.902	0.963	0.005	0.855	1.123
Profitability, serial correlation	0.818	0.755	0.006	0.733	0.773
Distance criterion (mean of Rdev2)			0.034		

Table reports the fit of the model to the statistics of the data that we target. See main text for details on the estimation method.

Table 2: Benchmark: Model Parameter Estimates

Parameter	Symbol	Estimate	SE
Preferences			
Risk Aversion	γ	105.856	8.931
Elasticity of intertemporal substitution	θ	2.207	5.216
Effective discount rate	ρ	0.040	0.025
Preference weight on relative consumption	h	0.947	0.231
Technology and Production			
Decreasing returns to investment	α	0.362	0.119
Depreciation rate	δ	0.033	0.260
Disembodied (complementary) technology growth, mean	μ_x	0.016	0.123
Disembodied (complementary) technology growth, vol	σ_x	0.077	0.009
Embodied (displacive) technology growth, mean	μ_ξ	0.004	0.242
Embodied (displacive) technology growth, vol	σ_ξ	0.137	0.012
Probability of becoming a low-growth firm	μ_L	0.283	0.343
Probability of becoming a high-growth firm	μ_H	0.015	0.029
Project arrival rate, low growth state	λ_L	0.122	0.018
Project arrival rate, high growth state	λ_H	8.588	2.047
Project-specific productivity, vol	σ_u	0.636	0.149
Project-specific productivity, persistence	κ_u	0.303	0.032
Incomplete Markets			
Fraction of project NPV that goes to inventors	η	0.785	0.376

Table reports the estimated parameters of the model. See main text for details on the estimation method.

Table 3: Comparison across restricted models

	DATA	BENCH	Restricted versions			Extension
			(R1) Complete Markets $[\eta = 0]$	(R2) No Relative Cons Pref $[h = 0]$	(R3) No Embod. Tech Shock $[(\mu_\xi, \sigma_\xi) = 0]$	(X1) LimPart $[q_S < 1]$
Aggregate quantities						
Consumption growth, mean	0.015	0.014	0.013	0.017	0.012	0.015
Consumption growth, volatility	0.036	0.042	0.042	0.036	0.040	0.038
Consumption growth, Long-Run risk	0.041	0.056	0.055	0.048	0.055	0.051
Shareholder Consumption, mean share of total	0.429	-	-	-	-	0.362
Shareholder Consumption, volatility	0.037	-	-	-	-	0.040
Investment, mean share of output	0.089	0.055	0.087	0.078	0.146	0.078
Investment, volatility	0.130	0.116	0.099	0.088	0.067	0.118
Investment and consumption, correlation	0.472	0.383	0.433	0.372	0.989	0.345
Net payout to assets, coefficient of variation	0.575	0.520	0.378	0.509	0.261	0.486
Asset Prices						
Market Portfolio, excess returns, mean	0.063	0.068	0.012	0.061	0.066	0.077
Market Portfolio, excess returns, std	0.185	0.152	0.133	0.122	0.095	0.134
Risk-free rate, mean	0.020	0.025	0.024	0.020	0.023	0.021
Risk-free rate, volatility	0.007	0.007	0.007	0.006	0.003	0.007
Value factor, mean	0.065	0.054	-0.004	0.004	-0.007	0.058
Value factor, volatility	0.243	0.191	0.100	0.063	0.236	0.168
Value factor, CAPM alpha	0.040	0.034	-0.004	0.001	0.018	0.037
Cross-sectional (firm) moments						
Investment rate, IQR	0.175	0.200	0.187	0.153	0.077	0.167
Investment rate, serial correlation	0.223	0.191	0.186	0.238	0.181	0.203
Investment rate, correlation with Tobin's Q	0.237	0.205	0.204	0.281	0.155	0.171
Tobin's Q, IQR	1.139	0.750	0.699	0.452	0.347	0.908
Tobin's Q, serial correlation	0.889	0.924	0.933	0.880	0.814	0.946
Profitability, IQR	0.902	0.963	0.930	0.827	0.808	0.927
Profitability, serial correlation	0.818	0.755	0.930	0.880	0.728	0.848
Distance criterion (mean of Rdev2)		0.034	0.192	0.151	0.255	0.026

Table compares the fit of the baseline model (column Bench) to three restricted versions of the model: a model where shareholders capture full rents from innovation, or $\eta = 0$ (column A); a model with complete markets that nests these two cases (column C); and a model where households have no preferences over relative consumption, or $h = 0$. See main text for details on the estimation method.

Table 4: Parameter Estimates: Comparison across restricted models

Parameter	Symbol	BENCH	Restricted versions			Extension
			(R1) Complete Markets	(R2) No Relative Cons Pref	(R3) No Embod. Tech Shock	(X1) LimPart $[q_S < 1]$
Preferences						
Risk Aversion	γ	105.856	94.682	17.826	89.860	58.250
Elasticity of intertemporal substitution	θ	2.207	2.088	1.305	1.947	1.851
Effective discount rate	ρ	0.040	0.028	0.045	0.059	0.046
Preference weight on relative consumption	h	0.947	0.990	-	0.849	0.804
Technology and Production						
Decreasing returns to investment	α	0.362	0.372	0.313	0.598	0.532
Depreciation rate	δ	0.033	0.029	0.056	0.087	0.029
Disembodied (complementary) technology growth, mean	μ_x	0.016	0.016	0.021	0.015	0.015
Disembodied (complementary) technology growth, vol	σ_x	0.077	0.078	0.064	0.080	0.070
Embodied (displacive) technology growth, mean	μ_ξ	0.004	0.000	0.002	-	0.006
Embodied (displacive) technology growth, vol	σ_ξ	0.137	0.127	0.106	-	0.113
Probability of becoming a low-growth firm	μ_L	0.283	0.452	0.277	0.351	0.353
Probability of becoming a high-growth firm	μ_H	0.015	0.009	0.059	0.056	0.025
Project arrival rate, low growth state	λ_L	0.122	0.152	1.821	0.235	0.111
Project arrival rate, high growth state	λ_H	8.588	6.708	3.642	17.744	11.983
Project-specific productivity, vol	σ_u	0.636	0.258	0.349	0.647	0.466
Project-specific productivity, persistence	κ_u	0.303	0.050	0.100	0.318	0.160
Incomplete Markets						
Fraction of project NPV that goes to inventors	η	0.785	0.000	0.727	0.900	0.598
Probability of born a stockholder	q_S	-	-	-	-	0.073

Table compares estimated parameters in the baseline model (column Bench) to four restricted versions of the model: a model where all households participate in financial markets, or $q_S = 1$ (column A); a model where shareholders capture full rents from innovation, or $\eta = 0$ (column B); a model with complete markets that nests these two cases (column C); and a model where households have no preferences over relative consumption, or $h = 0$. See main text for details on the estimation method.

Table 5: Consumption and Asset returns – Data versus Model

Asset	Variable	Consumption growth	Data	Model			
				Baseline	p(D=M)	LimitedPart	p(D=M)
Market	Returns	Aggregate	0.395	0.627	0.154	0.639	0.138
			[0.147]	[0.069]		[0.073]	
		Shareholders	0.212			0.717	0.002
			[0.117]			[0.118]	
		Shareholders,	0.111			0.296	0.477
		relative to non-shareholders	[0.154]			[0.21]	
Value Factor	Returns	Aggregate	0.162	-0.015	0.474	-0.009	0.458
			[0.13]	[0.209]		[0.19]	
		Shareholders	0.351			0.178	0.510
			[0.184]			[0.187]	
		Shareholders,	0.410			0.457	0.824
		relative to non-shareholders	[0.182]			[0.112]	

Table compares the empirical correlations between three consumption measures (aggregate consumption, consumption of stock holders, and relative consumption of stockholders) with the corresponding correlations in the model. The empirical correlations are based on the consumption data in (Malloy et al., 2009), covering the 1982-2002 period. The columns labeled $p(D = M)$ contain the p-values of the J-test of the hypothesis that the empirical correlations are equal to those in the model.

Table 6: Technology Shocks and Firm Profitability

A. Data							
Horizon	1	2	3	4	5	6	7
$\hat{\omega}_{I\setminus ft}$	-0.039	-0.040	-0.049	-0.049	-0.055	-0.060	-0.059
	(-11.24)	(-8.27)	(-7.85)	(-6.85)	(-6.44)	(-6.09)	(-5.37)
$\hat{\omega}_{I\setminus ft} \times G_{ft}$	0.022	0.016	0.018	0.017	0.017	0.018	0.020
	(6.71)	(3.32)	(3.04)	(2.43)	(2.08)	(1.95)	(1.94)
B. Model							
Horizon	1	2	3	4	5	6	7
$\hat{\omega}_{I\setminus ft}$	-0.023	-0.032	-0.039	-0.044	-0.048	-0.050	-0.052
	[-0.038, -0.009]	[-0.06, -0.007]	[-0.076, -0.007]	[-0.092, 0]	[-0.102, 0.006]	[-0.11, 0.012]	[-0.115, 0.015]
$\hat{\omega}_{I\setminus ft} \times G_{ft}$	0.010	0.022	0.032	0.042	0.050	0.057	0.063
	[0, 0.018]	[0.006, 0.041]	[0.006, 0.063]	[0.006, 0.081]	[0.006, 0.099]	[0.005, 0.115]	[0.008, 0.127]

The table summarizes the estimated coefficients a_0 and a_1 from the following specification

$$s_{ft+T} - s_{ft} = (a_0 + a_1\, G_{ft})\ \hat{\omega}_{I\setminus ft} + a_t + a_I + c_1\, G_{ft} + c_2\, \hat{N}_{ft} + c_l\, s_{ft} + c_k\, k_{ft} + e_{t+T}.$$

Here $s_{ft} \equiv \log S_{ft}$ represents log firm profitability at time t (In the model, firms face no ongoing costs of operation, hence profits and sales coincide. Examining sales instead of profits yields similar results.) We construct an industry-level equivalent of equation (39), denoted by $\hat{\omega}_{I\setminus f}$, that excludes firm f. We allow the relation between technological progress and future firm profitability to vary as a function of a firm's Tobin's Q, based on our discussion in Section 3.2. Specifically, we classify firms as either value ($G_{ft} = 0$) or growth ($G_{ft} = 1$) depending on whether their Tobin's Q falls below or above the industry median at time t. The variable $\hat{N}_{ft} \equiv \sum_{j \in P_{ft}} \hat{v}_{jt}/M_{ft}$ controls for innovation outcomes by firm f in time t – that is, the analogue of the dN_{ft} term in (35). Consistent with (35) we also include controls for lagged log profits y_{ft} and log size k_{ft}. In addition, we include time and industry dummies, and cluster the standard errors at the firm level. We also estimate equation (5) in simulated data from the model. Since the model contains no industries, we replace $\hat{\omega}_{I\setminus f}$ with $\hat{\omega}$ and drop the time and industry dummies. To facilitate comparison between the data and the model, we scale $\hat{\omega}_{I\setminus f}$ to unit standard deviation.

Table 7: Measuring technological progress – Data versus model

Statistic	Data	Model			
		Bench	$p(D = M)$	LimitedPart	$p(D = M)$
$\hat{\omega}$, serial correlation	0.765	0.906	0.106	0.907	0.098
	[0.068]	[0.055]		[0.053]	
$\Delta\hat{\omega}$, volatility	0.303	0.156	0.003	0.152	0.002
	[0.047]	[0.012]		[0.015]	
$\Delta\hat{\omega}$, correlation with value factor	-0.284	-0.350	0.606	-0.353	0.621
	[0.078]	[0.102]		[0.115]	
$\Delta\hat{\omega}$, correlation with market portfolio	-0.499	-0.419	0.474	-0.379	0.317
	[0.088]	[0.067]		[0.081]	
$\Delta\hat{\omega}$, correlation with agg. cons. growth	-0.216	-0.001	0.250	0.003	0.249
	[0.113]	[0.148]		[0.153]	
$\Delta\hat{\omega}$, correlation with shareholder abs. cons. growth	-0.303			-0.229	0.788
	[0.137]			[0.235]	
$\Delta\hat{\omega}$, correlation with shareholder rel. cons. growth	-0.354			-0.643	0.081
	[0.126]			[0.108]	

Table the moments of the measure of the value of new projects to the value of the market portfolio, $\hat{\omega}$, in the data and in the model. See main text for details on the construction of $\hat{\omega}$. Numbers in brackets are standard errors.

Table 8: Estimates of a linearized Stochastic Discount Factor (Data vs Model)

	A. Data						B. Model					
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(i)	(ii)	(iii)	(iv)	(v)	(vi)
Δx	1.89	0.41					0.50	0.28				
	[0.85]	[0.76]					[0.18]	[0.21]				
$\Delta \log C$			1.79	0.01					0.80	0.46		
			[0.95]	[0.35]					[0.31]	[0.32]		
$\Delta \log Y$					2.06	0.37					0.83	0.45
					[1.2]	[0.35]					[0.33]	[0.32]
$\Delta\hat{\omega}$		-0.80		-0.83		-0.69		-0.69		-0.63		-0.69
		[0.38]		[0.27]		[0.24]		[0.38]		[0.39]		[0.39]
R2	0.35	0.86	-0.20	0.84	-1.38	0.80	-0.29	0.50	-0.20	0.50	-0.48	0.50
	[3.12]	[0.22]	[3.63]	[0.1]	[3.53]	[0.11]	[0.44]	[0.4]	[0.57]	[0.38]	[0.74]	[0.38]
MAPE (%)	1.28	0.66	1.93	0.72	2.89	0.86	1.68	0.93	1.60	0.92	1.76	0.92
	[1.13]	[0.32]	[1.69]	[0.28]	[1.67]	[0.25]	[0.53]	[0.31]	[0.53]	[0.31]	[0.58]	[0.31]
J-stat	21.71	9.05	39.55	9.99	41.08	9.99	39.75	26.23	40.45	26.54	42.30	26.54
	[21.35]	[14.3]	[13.74]	[8.79]	[14.06]	[9.03]	[24.23]	[20.71]	[26.04]	[21.11]	[26.15]	[21.1]
(p-val)	0.01	0.43	0.00	0.35	0.00	0.35	0.02	0.13	0.02	0.13	0.02	0.13

Table shows estimates of a linearized SDF $m = a - bF$ using the cross-section of 10 portfolios sorted on book-to-market. We normalize a so that $E[m] = 1$, which yields the moment conditions $E[R^e] = -cov(R^e, m)$. We consider various combinations of F, including utilization-adjusted TFP (proxied by Δx in the model), aggregate consumption growth, aggregate output (GDP) growth and changes in $\hat{\omega}$. Standard errors in the data are HAC (3-lags) and include a Shanken adjustment. We report mean parameter estimates across simulations and 90% confidence intervals. In addition, we report the cross-sectional R^2, the mean absolute pricing error (MAPE), and the Hansen J statistic. Standard errors for the last three statistics are computed using a delete-d Jackknife estimator, where $d = 2\sqrt{T}$.

Table 9: Technology shocks and consumption growth

	Horizon (years, T)				
	(1)	(2)	(3)	(4)	(5)
Data (no hh controls)					
$\hat{\omega}_T - \hat{\omega}_t$	-0.209	-0.480	-0.413	-0.379	-0.318
	[0.26]	[0.21]	[0.17]	[0.18]	[0.25]
Data (hh controls)					
$\hat{\omega}_T - \hat{\omega}_t$	-0.348	-0.465	-0.392	-0.470	-0.707
	[0.21]	[0.15]	[0.19]	[0.31]	[0.34]
MODEL					
$\hat{\omega}_T - \hat{\omega}_t$	-0.168	-0.223	-0.266	-0.302	-0.332
	[0.12]	[0.16]	[0.18]	[0.20]	[0.21]

Table shows estimates b_1 from the following linear regression

$$(c_T^{md} - c_T) - (c_t^{md} - c_t) = b_1\,(\hat{\omega}_T - \hat{\omega}_t) + b_2\,(c_t^{md} - c_t) + b_3\,\hat{\omega}_t + u_t \tag{5}$$

where c is aggregate log per capita consumption expenditures, c^{md} is median log per capital consumption (from the CEX). We construct using total expenditures (first row) and also as residuals from household level controls (family size, number of earners, member size, years of schooling and a quadratic age term). Data period is 1982-2010. Simulated samples are same length. Standard errors (in brackets) are HAC, with maximum lag length equal to $T + 1$. All variables are standardized to mean zero and unit standard deviation.

Figure 2: Technology and Aggregate Quantities

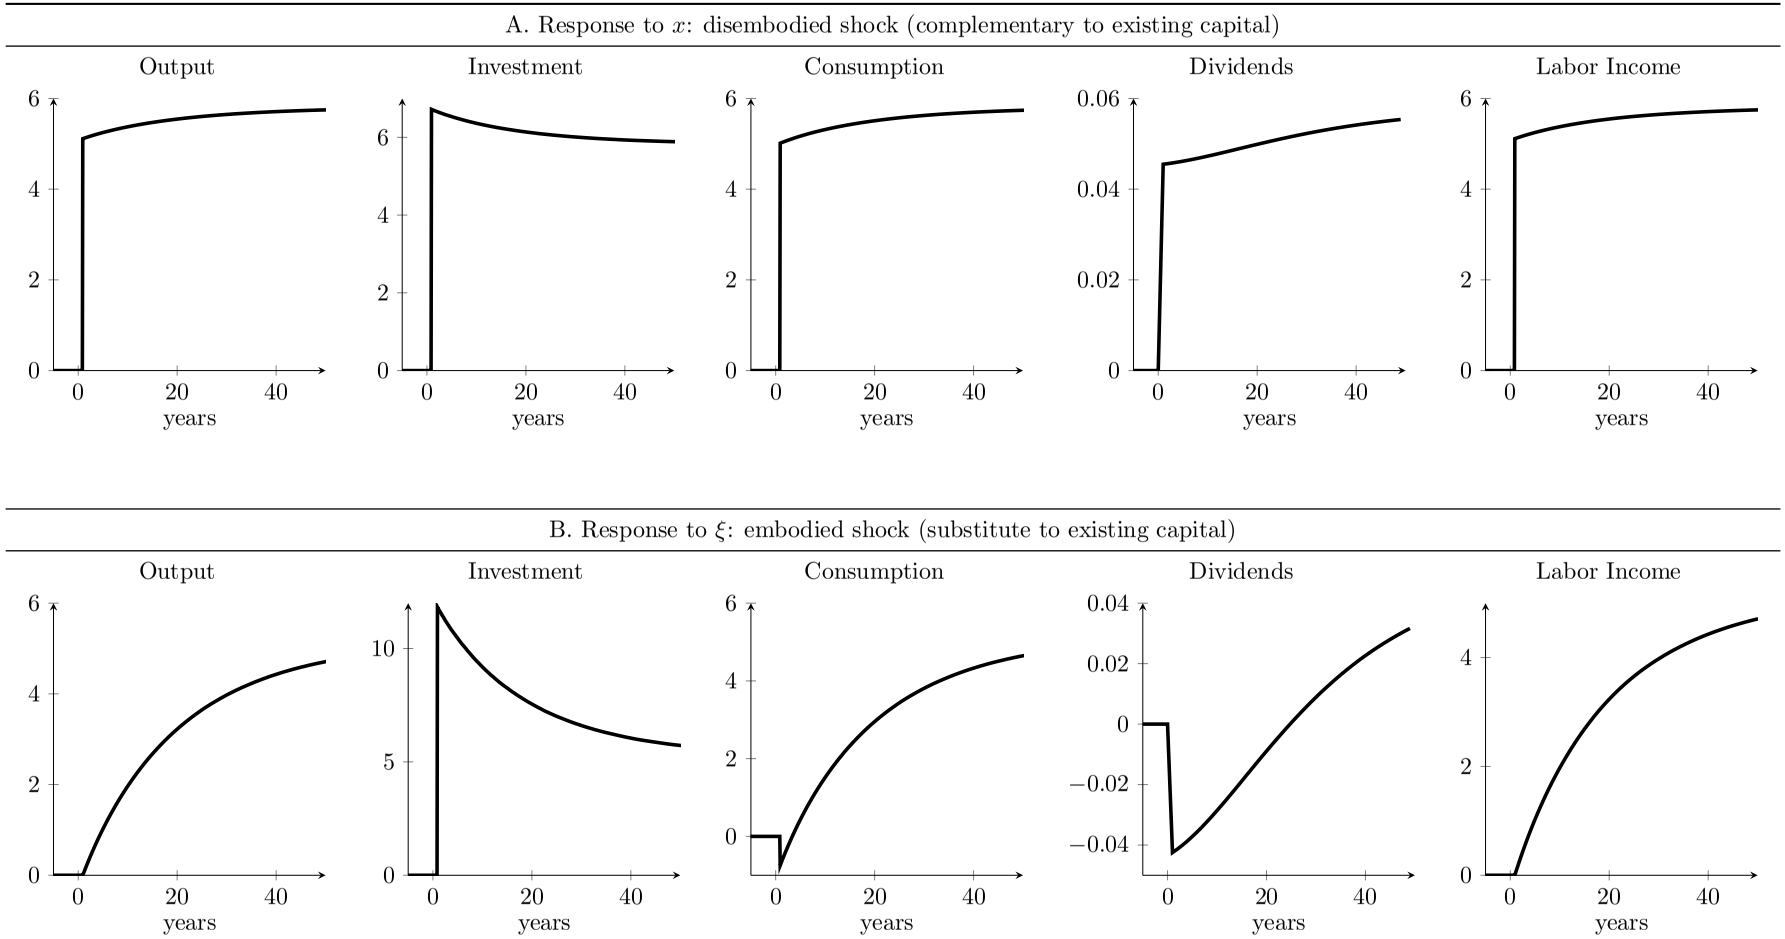

Figure plots the impulse response of aggregate output, investment and consumption expenditures to the two technology shocks in the model. We construct the impulse responses taking into account the nonlinear nature of equilibrium dynamics: we introduce an additional one-standard deviation shock at time t without altering the realizations of all future shocks.

Figure 3: Technology and Asset Prices

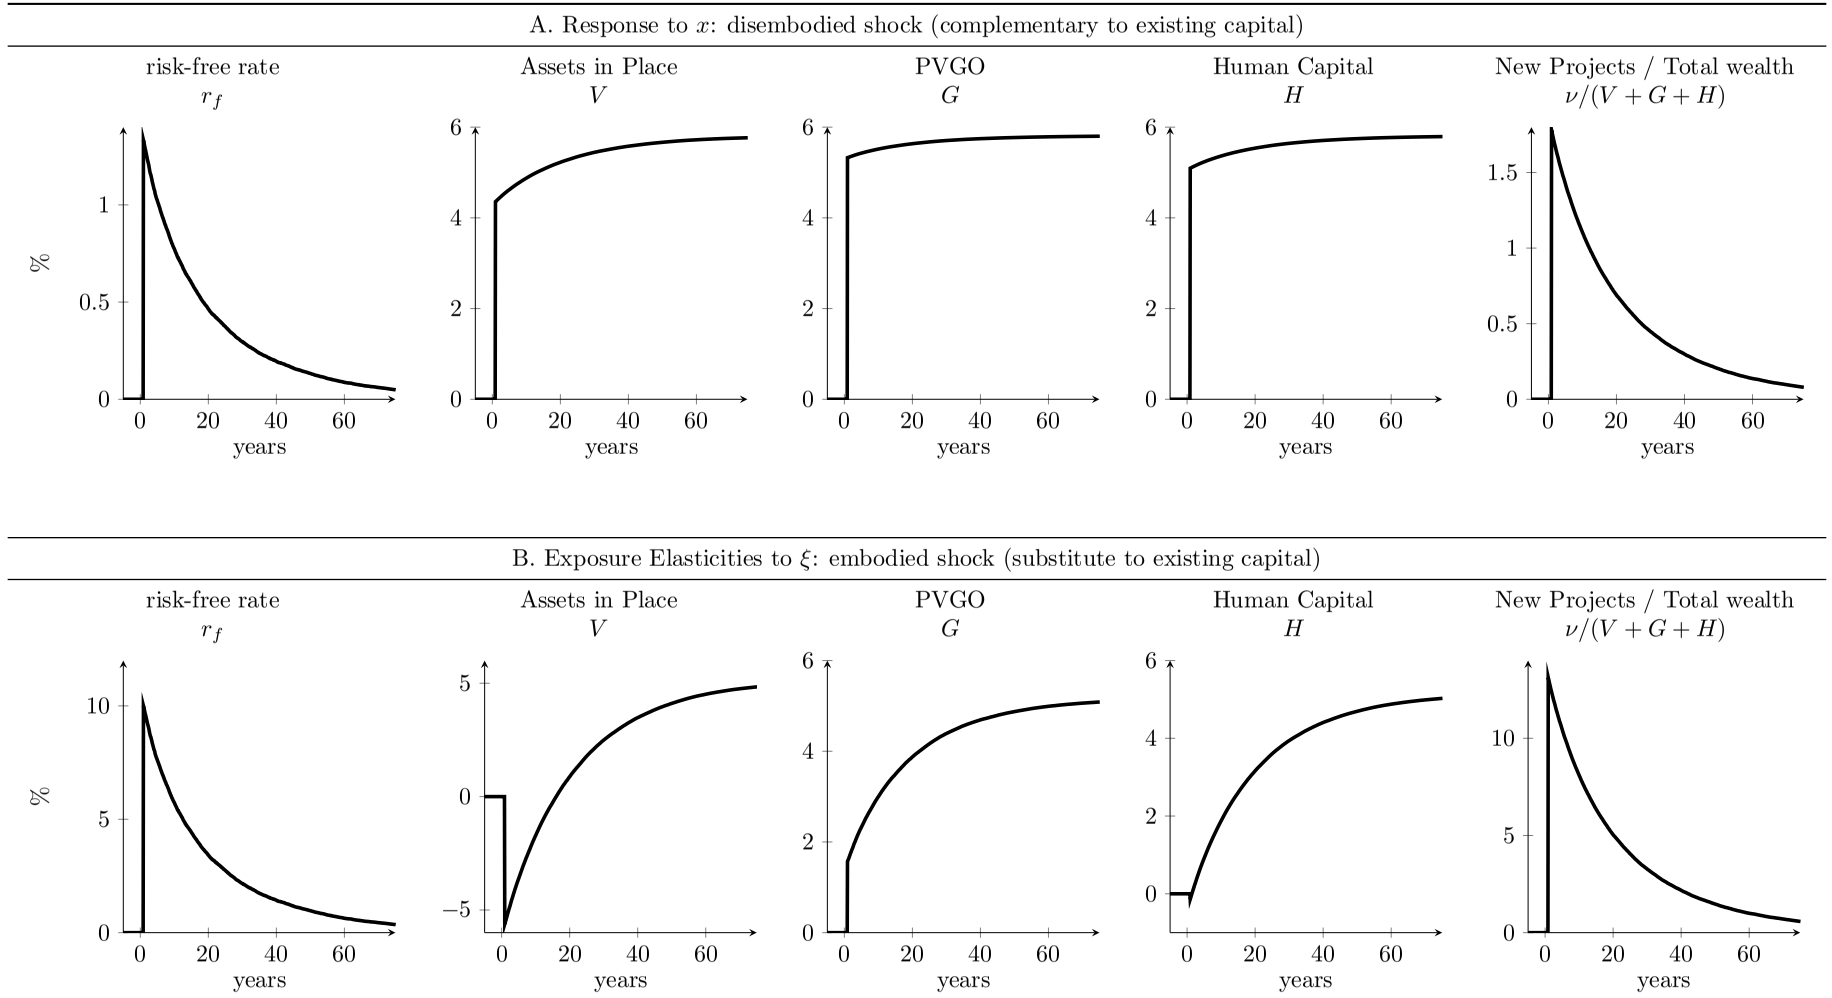

Figure plots the impulse response of aggregate output, investment and consumption expenditures to the two technology shocks in the model. We construct the impulse responses taking into account the nonlinear nature of equilibrium dynamics: we introduce an additional one-standard deviation shock at time t without altering the realizations of all future shocks.

Figure 4: Technology and Consumption of Current Market Participants

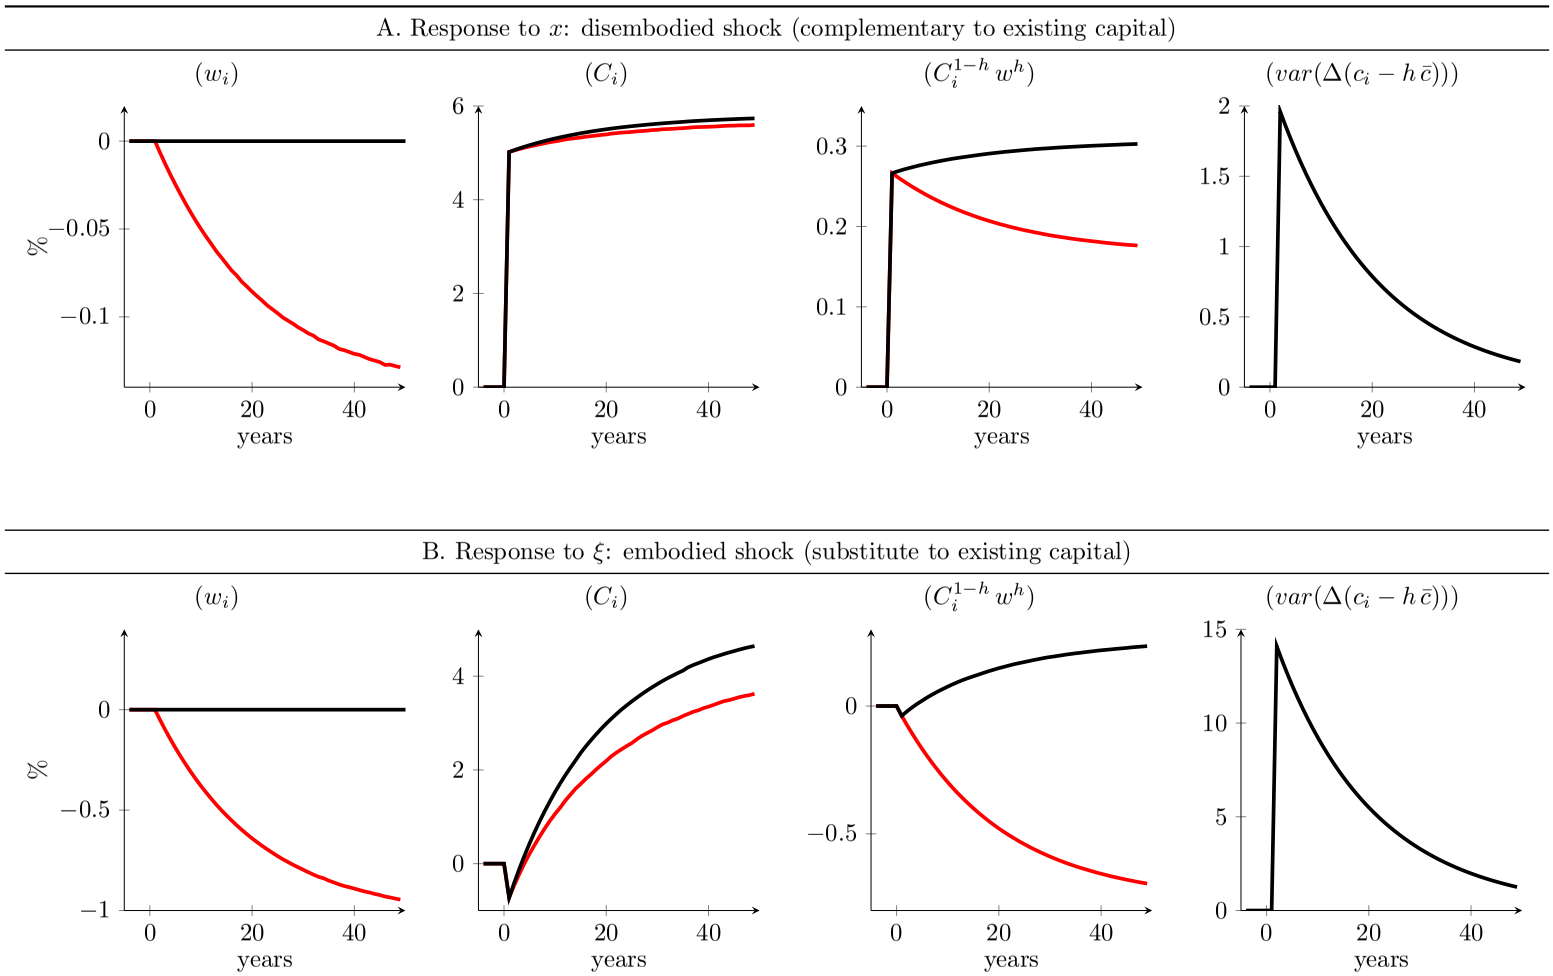

The first three columns plot the impulse response of the household wealth share (w_i), household consumption (C_i), and consumption adjusted for relative consumption preferences $C_i^{1-h}\, w_i^h$. We construct the impulse responses taking into account the nonlinear nature of equilibrium dynamics: we introduce an additional one-standard deviation shock at time t without altering the realizations of all future shocks. We plot the response of the mean in black, and the response of the median in red. The last row plots the impulse response of the conditional variance of instantaneous log consumption growth, adjusted for relative preferences.

Figure 5: Technology and the Stochastic Discount Factor

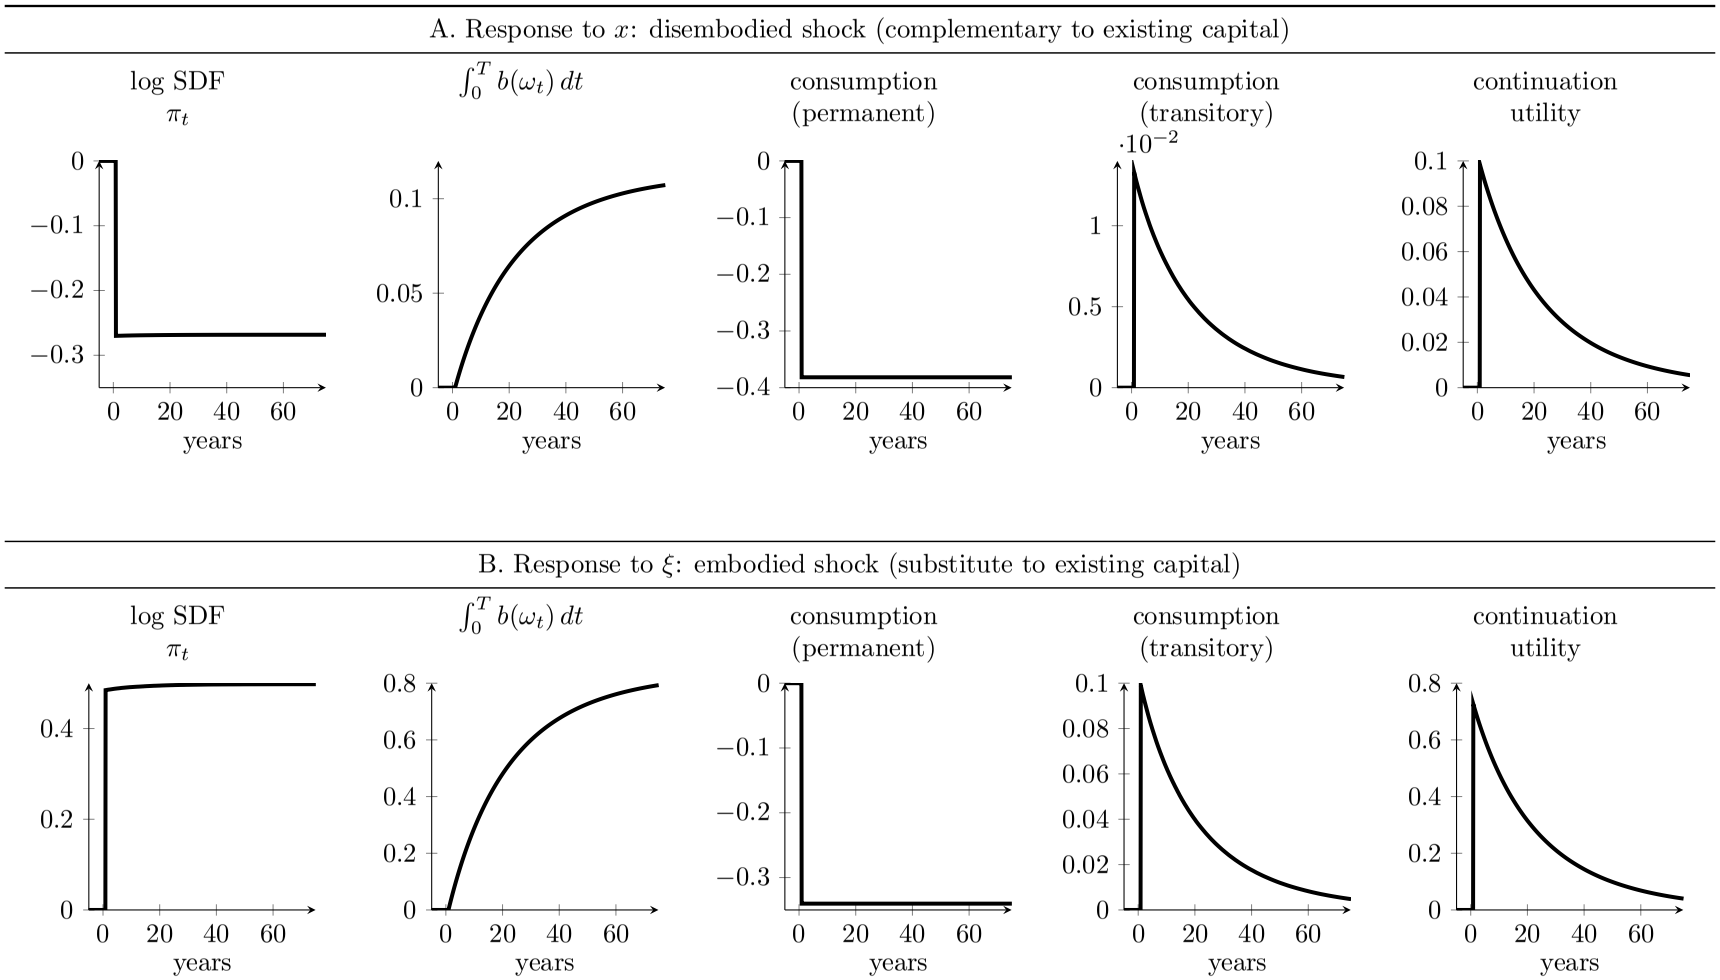

Figure plots the impulse response of the log stochastic discount factor – and its four components – to the two technology shocks in the model. We construct the impulse responses taking into account the nonlinear nature of equilibrium dynamics: we introduce an additional one-standard deviation shock at time t without altering the realizations of all future shocks.

Figure 6: Technological shocks and Firms

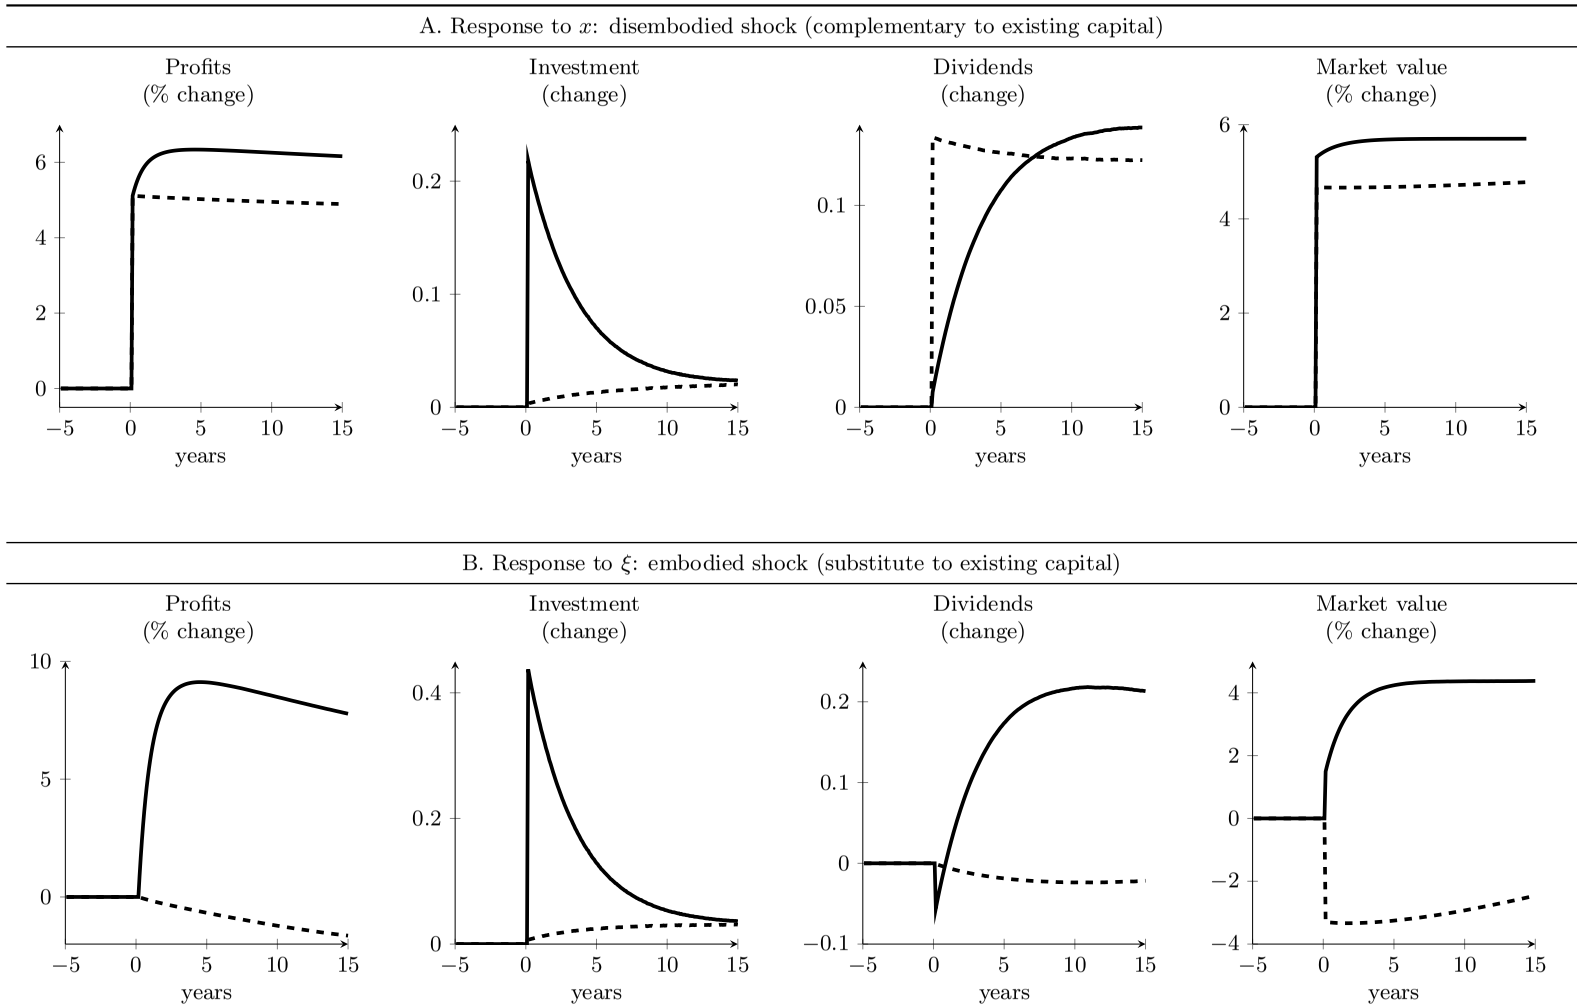

Figure plots the dynamic response of firm profits, investment, dividends and stock prices to the two technology shocks x and ξ in the model. We report separate results for two types of firms. The solid line represents the responses for a Growth firm, defined as a firm with $\lambda_{ft} = \lambda_H$ and $k_{ft} = 0.5$. The dotted line indicates the responses for a value firm, defined as a firm with low investment opportunities $\lambda_{ft} = \lambda_L$ and large size $k_{ft} = 2$. For both firms, the level of average profitability is equal to its long-run mean, $u_{ft} = 1$. The initial value of the state variable ω is set to its unconditional mean, $\omega_0 = \mathrm{E}[\omega_t]$. Columns 1 and 4 plot percentage changes, columns 2 and 3 plot changes in the level (since both dividends and investment need not be positive).

Figure 7: Firm risk exposures and risk premia

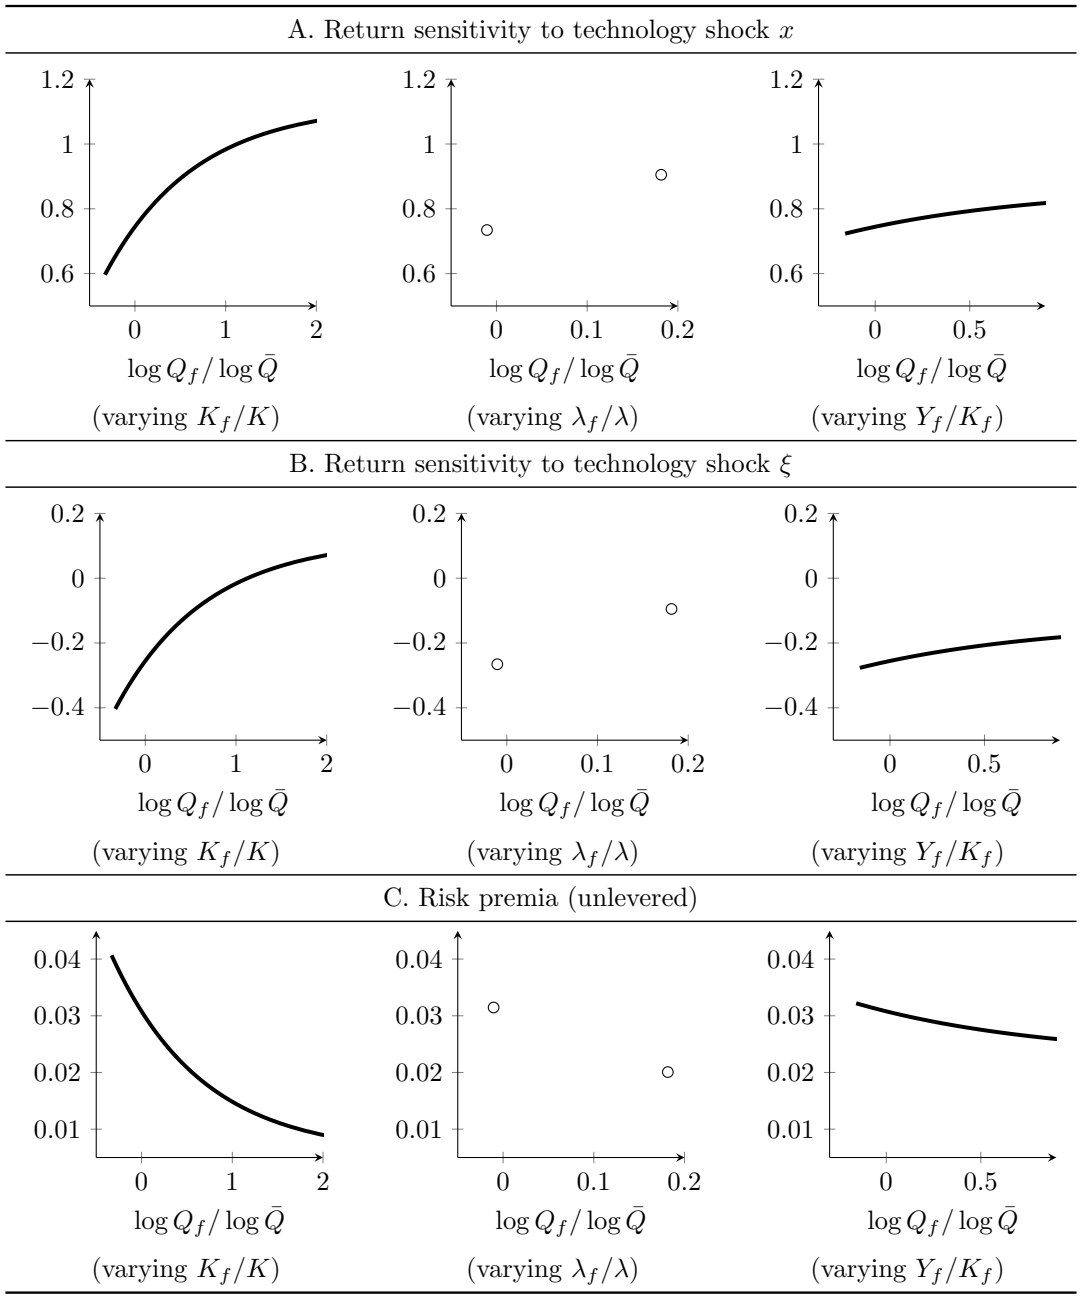

Figure illustrates how technology risk exposures (Panels A and B) and risk premia (Panel C) vary with the firm's market-to-book ratio (Q). A firm's market to book ratio is a function of the firm's relative size K_f/K, likelihood of future growth λ_f/λ, and its current profitability Y_f/K_f. In each of the three columns, we examine how variation in Q due to each of these three state variables translates into variation in risk premia. The range in the x-axis corresponds to the 0.5% and 99.5% of the range of each these three variables in simulated data.

Figure 8: Estimated value of new technologies

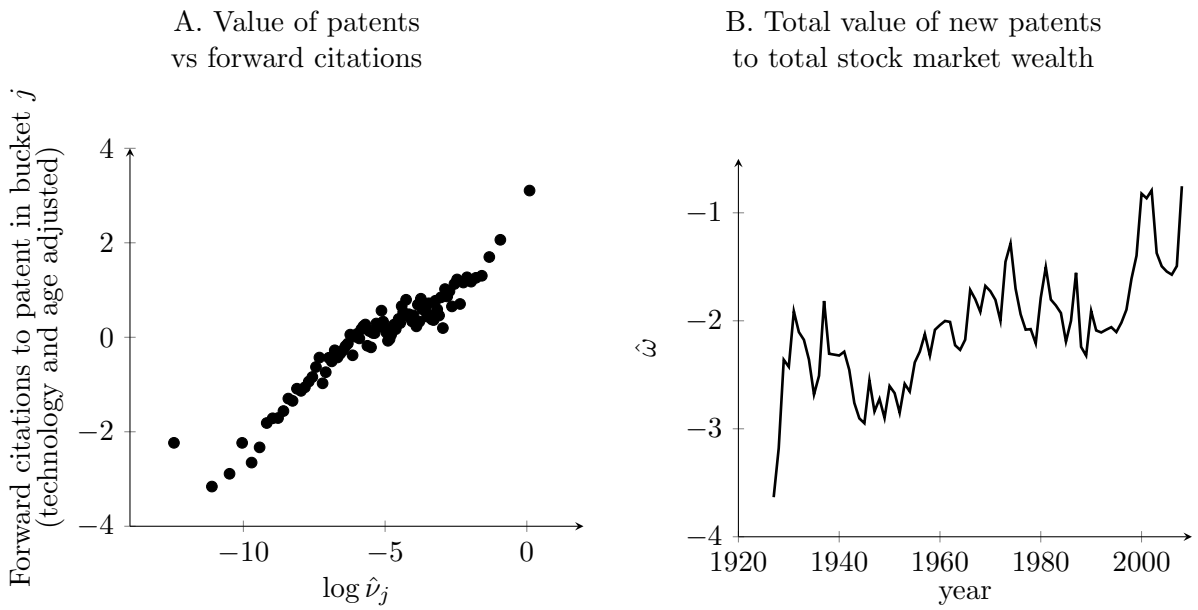

Panel A plots the relation between forward patent citations and the estimated market value of patents. We group the patent data into 100 quantiles. The vertical axis plots the average number of patent citations in each quantile–minus the average number of citations to patents in the same technology class that were granted in the same year. The horizontal axis plots the logarithm of the average estimated patent value in each quantile. Panel B plots the time series of the ratio of the estimated value of new technologies $\hat{\nu}_t$ to the value of the stock market. See the main text, the appendix and Kogan et al. (2012) for more details.