

Legacy Events Room
CBA 3.202
Thursday, March 12, 2020
11:00 am

Finance in a Time of Disruptive Growth

Nicolae Gârleanu

UC Berkeley-Haas and NBER

Stavros Panageas

UCLA, Anderson School of Management and NBER

November 2019

Incomplete Draft

Abstract

We build a model in which the arrival of new technologies displaces demand for old technologies. This disruption causes redistribution due to lack of risk sharing both within and across investor cohorts. We model the financial industry as a costly device to improve risk sharing, and determine its size in equilibrium. We further study wealth dynamics, equilibrium prices, and flows into various asset classes. We show that an increase in disruptive activity renders existing firms' publicly traded equities riskier and renders "alternative asset classes" as diverse as fixed income, real estate, and private equity more attractive. The result is a decline in the real interest rate, an expansion of the financial industry, and increased flows towards alternative asset classes. Interestingly, alternative asset classes offer higher expected rates of return than conventional equities despite the diversification benefits afforded by the former.

1 Introduction

Two salient trends in the financial markets over the past few decades have been the protracted decline of real interest rates and the increased portfolio allocations to alternative classes — asset families as different as private equity, venture capital, and commercial real estate. We link these trends to an increased incidence of “disruptive” — or, more appropriately, redistributive — growth. Specifically, growth leads new firms to capture a larger portion of profits and market capitalization — at the expense of old firms, which get displaced. Crucial, the creation of new firms benefits investors asymmetrically: the benefits accrue predominantly to the firm creators (and, more generally, early investors holding large fractions of their equity), rather than to investors simply holding the market portfolio of public companies.

If the dispersion of wealth growth across investors increases, then investors have an incentive to increase allocations towards asset classes that are less affected by, or can even benefit from, disruption risk. Our model identifies private equity, real assets, and risk-free assets as asset classes that fit that requirement, thus explaining their popularity in recent decades. Somewhat surprisingly, the equilibrium expected rate of return of funds investing in emerging technologies exceeds that of the market portfolio of existing equities, despite the positive market price of risk for displacement risk and the negative exposure of these emerging technologies to displacement risk.

We develop a general equilibrium model with the following features. The production of a final good requires labor and intermediate products. New lines of (intermediate) products, which are introduced in stochastic amounts, raise aggregate production, but also displace the demand for, and hence the profits generated by, old intermediate products. The ownership rights to the production of the new product lines are allocated either to existing publicly traded firms or to newly arriving agents. The allocation of blueprints to new agents is random, and highly asymmetric. A small number of these agents end up with the profitable product lines, while the rest receive worthless allocations. As a result, the newly arriving agents are eager to share the allocation risk with investors, by offering a fraction of their firm’s shares for sale. This transaction is facilitated by financial intermediaries who purchase a portfolio of the new-firm shares on behalf of investors. This diversification, however, is costly; in particular, each intermediary optimally invests only in a subset of new firms.

There are therefore two impediments to perfect risk sharing. At the inter-cohort level, the fraction retained by the newly arriving agents is excessive. At the intra-cohort level, the portfolio of new firms available to a typical investor, via an intermediary, is imperfectly

correlated with the aggregate value of all new — non-public — firms. The model nests the perfect risk-sharing limit (Rubinstein (1976), Lucas (1978)), the Constantinides and Duffie (1996) model, and the OLG model of Gârleanu et al. (2012) as special parametric cases.

We provide explicit closed-form solutions of the model and show the following results. First, if risk sharing is close to perfect, then an increased arrival of new technologies is “good news” for the marginal investor, since new technologies are good news for aggregate output and market capitalization. However, if either intra- or inter-cohort risk sharing fails, then increased arrival of new technologies is perceived as a negative outcome by the representative investor, who might end up losing from the new technologies.

Second, a surprising result of the model is that the equilibrium returns of emerging, privately held firms must exceed the expected returns of publicly traded firms, even though they offer hedging opportunities by reducing the exposure of an investor’s portfolio to displacement risk. The reason is that in a world of imperfect risk sharing the diversifiable risk of private investments commands risk compensation.

Third, an acceleration in displacement or dispersion of innovation gains across investors increases the incentives to diversify out of public equities. The natural targets are risk-free assets, private equities, and real assets. Real assets, such as real estate and commodities, benefit from increased arrival of new firms because they are useful to all firms (new and old). Private equity benefits because it helps offset the displacement risk of public equities. Finally, the demand for the risk-free asset increases due to precautionary savings incentives, leading to a decline in the risk free rate.

Fourth, an acceleration in displacement leads to an increase in the size of the financial industry and a reduction in real interest rates, though possibly not to increased physical investment. The reason is that, if investment is specific to blueprints, then increased displacement raises the risk that these investments render themselves unprofitable. This may help explain why the low real rates observed in recent decades did not lead to a substantial increase in investment, but rather an expansion of the financial industry.

In short, the paper makes two contributions. First, it introduces multiple asset classes, distinguished by their different relations to displacement risk, in a general equilibrium model and provides a rationalization for the diverse average returns that these asset classes have offered historically. Second, in a world of increased disruptive activity, the model provides a possible unified rationalization for several patterns observed in recent decades, and especially the simultaneous increase in demand for risk-free investments and relatively risky alternative asset classes.

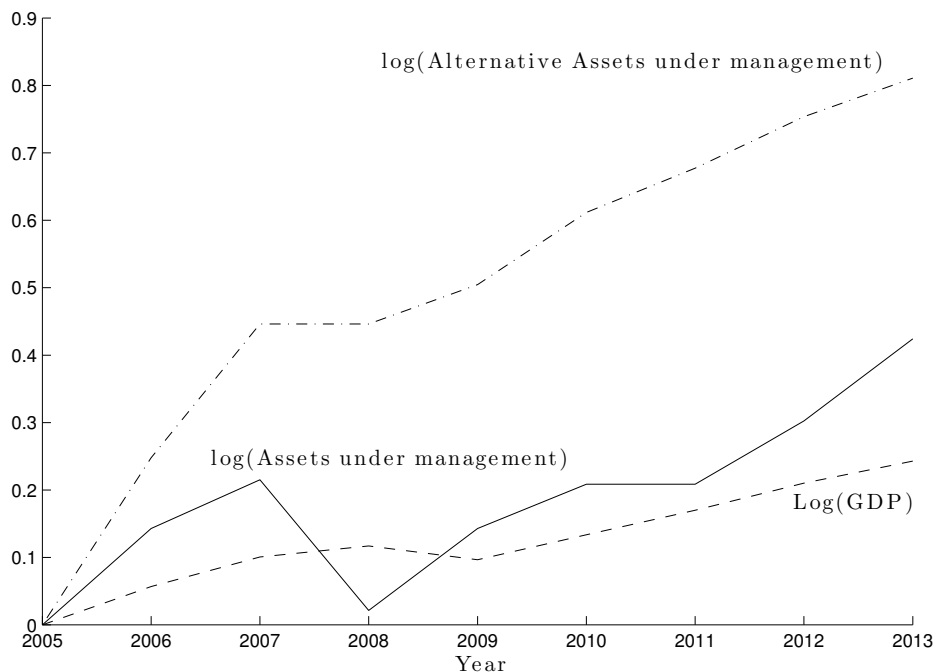


Figure 1: Cumulative first differences in the logarithm of GDP, assets managed by the financial industry, and alternative assets in the form of private equity, venture capital, and commercial real estate. Source: McKinsey (2015).

1.1 Literature Review

To be written.

2 Empirical motivation

Figure 1 illustrates the growth in alternative assets over the last decade. An obvious conclusion is that the alternative asset management industry in the form of private equity, venture capital, and commercial real estate has grown much faster than either GDP or total assets under management in the economy. This substantial growth in the size of alternative investments coincided with substantial growth in other forms of alternative investments (such as commodities or hedge funds) and with a protracted period of declining real interest rates, trends that have been documented repeatedly in the literature.

Our goal in this paper is to explain these trends not as isolated phenomena, but rather as emanating from a common source, namely the increased incidence of “disruptive growth,”

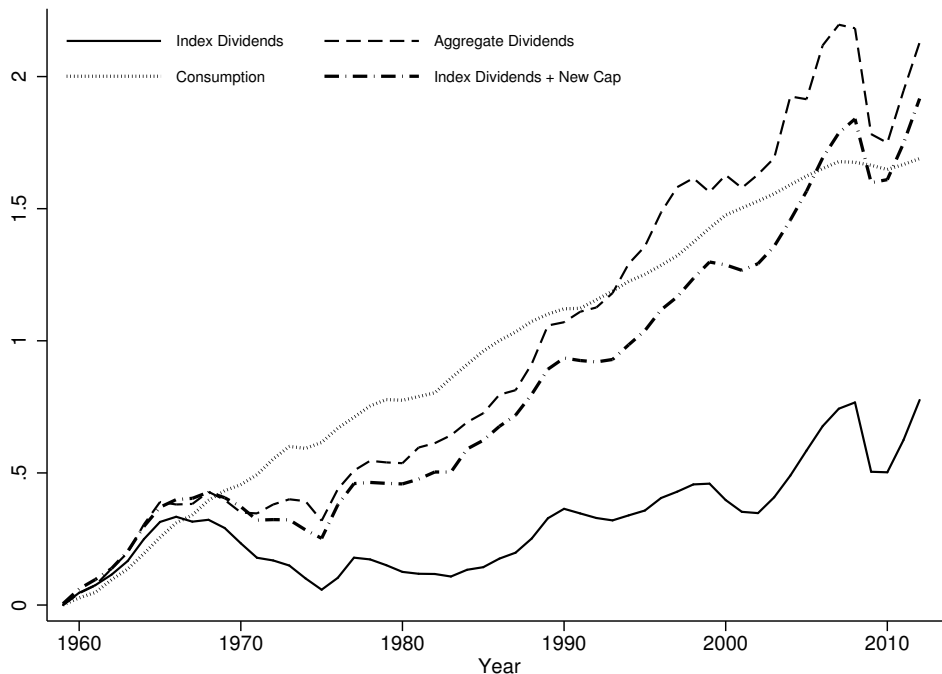


Figure 2: Total real logarithm of S&P 500 dividends per share, real log-aggregate consumption and real log-aggregate dividends. The CPI is used as a deflator for all series. The line “Index Dividends + New Cap” is equal to real log-dividends per share plus the cumulative (log) gross growth in the shares of the index that are due to the addition of new firms. Sources: R. Shiller’s website, FRED, Personal Dividend Income series, and CRSPSift.

or more accurately, redistributive growth.

To motivate these notions, we point to Figure 2. This figure shows that even though aggregate dividends and aggregate consumption share a common trend, the dividends per share of the S&P 500 follow a markedly slower growth path. (The same conclusion holds for the CRSP value weighted portfolio). The difference in growth rates between aggregate dividends and dividends per share is approximately 2% per year. This discrepancy can be largely attributed to a dilution effect arising every time new companies enter the index: Just as a corporation has to issue new shares to purchase a new company, thus diluting the existing shareholders, the same happens at the level of an index. Mathematically, if S_t is the divisor (of number of index “shares”) at time t , MV_{t+1}^{old} the market capitalization at time $t + 1$ of firms that were already part of the index at time t , and MV_{t+1}^{new} the market capitalization of firms that are added to the index at time $t + 1$, the insensitivity of the index

level to additions is written as

$$\frac{MV_{t+1}}{S_{t+1}} = \frac{MV_{t+1}^{\text{old}} + MV_{t+1}^{\text{new}}}{S_{t+1}} = \frac{MV_{t+1}^{\text{old}}}{S_t}$$

giving

$$\frac{S_{t+1}}{S_t} = 1 + \frac{MV_{t+1}^{\text{new}}}{MV_{t+1}^{\text{old}}}. \tag{1}$$

Next, we let D_t denote total dividends, and define $\frac{D_t}{S_t}$ as dividends per share. The concept of dividends per share corresponds to the cash flows generated by the self-financing strategy of holding the market (or index) portfolio. Being the cash flows of a self-financing strategy is helpful from an asset pricing perspective, since the standard arguments imply that the value of the index equals to the present value of the stream $\frac{D_t}{S_t}$.

Empirically, over long horizons the discrepancy of the trends between dividends-per-share and aggregate consumption is almost exclusively due to the diluting effect of new company additions. Figure 2 shows that if we add back $\sum_t \log(1 + MV_{t+1}^{\text{new}}/MV_{t+1}^{\text{old}})$ to the time series of (log) dividends-per-share $\log(D_t/S_t)$, then the resulting series (denoted “Index Dividends + New Cap”) co-trends with aggregate dividends and aggregate consumption. Hence the dillutionary effect of new company additions is the main reason for the discrepancy between the cash flows generated by the self-financing strategy of holding the market portfolio and aggregate consumption — a fact that our model reproduces.

Figure 3 shows that the dillutionary effect of new company additions has accelerated over time. The figure plots 20-year moving averages of $\log(1 + MV_{t+1}^{\text{new}}/MV_{t+1}^{\text{old}})$ and documents an upward trend.

Casual empiricism suggests that the benefits of the introduction of new firms are highly asymmetric. It is typically the firm creators and the early investors who benefit. The data offer some evidence in that direction. Figure 4 compares the total wealth in the hands of billionaires that inherited their wealth with the wealth of “self-made” entrepreneurs. In 2003 (the earliest electronically available cross section in Forbes) the estimated net worth of self made billionaires is roughly comparable to that of billionaires who inherited their wealth. By 2016, the net worth in the hands of self-made billionaires is roughly twice as large as the wealth of billionaires who inherited their wealth.

In addition, Table 1 shows that a billionaire who inherited her wealth has about the same wealth as a self-made billionaire in 2016. Moreover, for ages between 45 and 65 a self-made billionaire tends to have a higher net worth than someone who inherited their wealth. This didn’t use to be the case. In 2003, a self-made billionaire had on average about 31% less

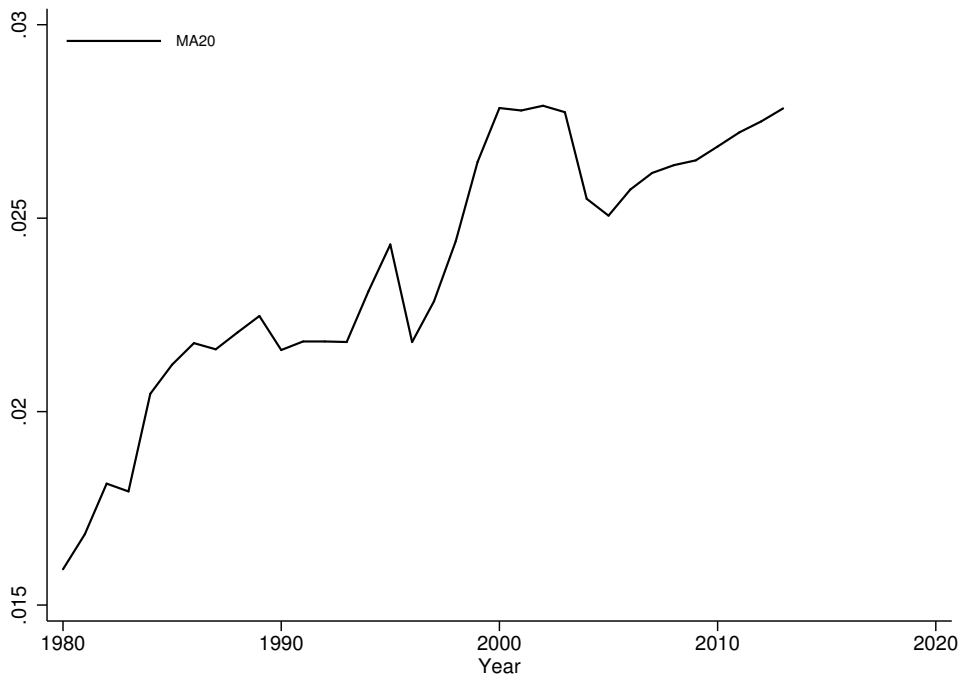


Figure 3: 20-year moving average of the change in the index divisor that is due to the addition of new firms.

wealth than a billionaire who inherited her wealth, with a t-stat of about 1.9. The difference in the 2003 cross section is even larger for older billionaires. Figure 6 shows that it is not only the averages that differ: in the 2016 cross-section the distribution of billionaire wealth looks identical for the subgroups of self-made and heirs. Taken together, Table 1 and Figure 6 imply that self-made billionaires must have experienced a stronger wealth growth (compared to heirs) in recent times.¹

Figure 5 performs a related exercise: The figure fixes billionaire families in the 2003 edition of Forbes. Using a newly available database called Wealth-X, we follow the wealth growth of each family between 2003 and 2016. This database is ideally suited for this purpose: it is professionally maintained by 170 employees who collect data pertaining to an individual’s publicly disclosed transactions, holdings, philanthropy, large purchases, board

¹To make this statement mathematically precise, express a billionaire’s wealth growth as $\log(W_T) = \log(W_0) + u$ where u captures the cumulative rate of growth in the value of a billionaire’s assets net of her consumption-to-wealth ratio. Figure 6 shows that the distribution of $\log(W_T)$ (conditional on W_T being above a billion) is independent of $\log(W_0)$. Hence, the distribution of u for self-made and non self-made billionaires cannot be the same. Indeed, the distribution of u for self-made billionaires must contain a positive location shift compared to the one for billionaires who inherited their wealth.

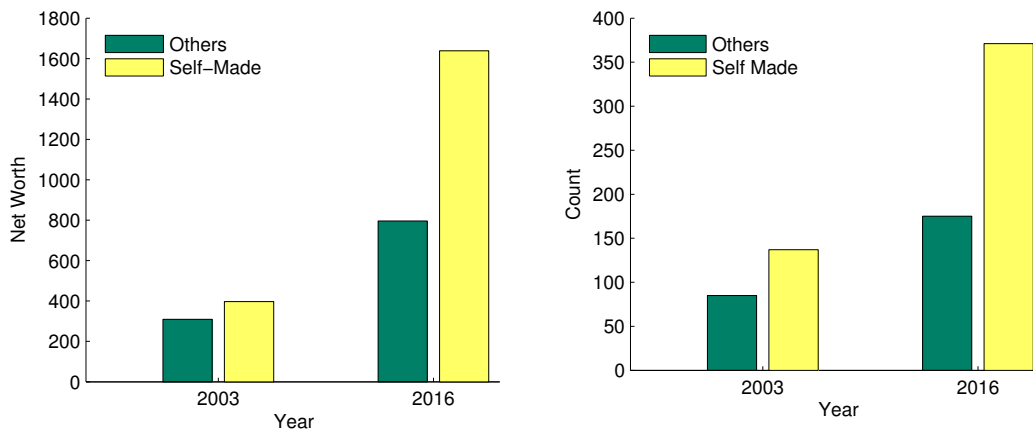


Figure 4: Wealth of US billionaires who inherited their wealth and those who were self-made (left figure). Number of US billionaires who inherited their wealth and number of billionaires who are self made. Source: Forbes 2003, 2016.

memberships, professional and family ties, etc., and aggregate them into a detailed “folder” for that individual and their family. In constructing family wealth growth we compare the wealth of each family in 2003 to the wealth of the same family in 2016.² The somewhat surprising fact about Figure 5 is that the median family experienced a wealth growth rate no different than the growth in the price-per-share of the S&P 500. Taken as group, the wealth of these families has about doubled between 2003 and 2016, while the wealth in the hands of billionaires has quadrupled between 2003 and 2016. This implies that the large increase in the hands of billionaires has been the result of entry of new wealth in that population, rather than the organic growth of the wealth of the previously rich.

Table 2 shows that the above results are unlikely to be the result of error or different consumption-to-wealth rates between new and old rich. Using data from the Wealth-X database, we regress an individual’s publicly known log philanthropic expenditure over her life-time on her estimated log-net worth and a dummy variable taking the value one if the billionaire is listed as “self-made.” This table shows two things: First the coefficient on log net worth is essentially equal to one. Since it is reasonable to expect that philanthropy should be proportional to an individual’s true, unobserved net worth, this is reassuring. It suggests that net worth is measured reasonably well, and does not just capture “paper money”; instead we see it reflected in an easily observed expenditure component. The second

²We include immediate descendants and divorced spouses as part of a family for the computation of total wealth growth rates.

Source: Wealth-X 2016			
	(1) All Ages	(2) 45 < Age < 65	(3) Age > 65
Self-Made Dummy	-0.029 (0.087)	0.049 (0.109)	-0.150 (0.155)
Number of Observations	413	165	215
R^2	0.000	0.001	0.006
Source:Forbes 2016			
	(1) All Ages	(2) 45 < Age < 65	(3) Age > 65
Self-Made Dummy	-0.010 (0.095)	0.124 (0.110)	-0.148 (0.177)
Number of Observations	465	197	233
R^2	0.000	0.005	0.004
Source: Forbes 2003			
	(1) All Ages	(2) 45 < Age < 65	(3) Age > 65
Self-Made Dummy	-0.310 (0.165)	-0.324 (0.240)	-0.492 (0.264)
Number of Observations	176	83	73
R^2	0.027	0.028	0.062

Table 1: Regressions of billionaire’s log(net worth) on a dummy variable taking the value one if the billionaire is characterized as self made in the respective data set. Standard errors in parentheses

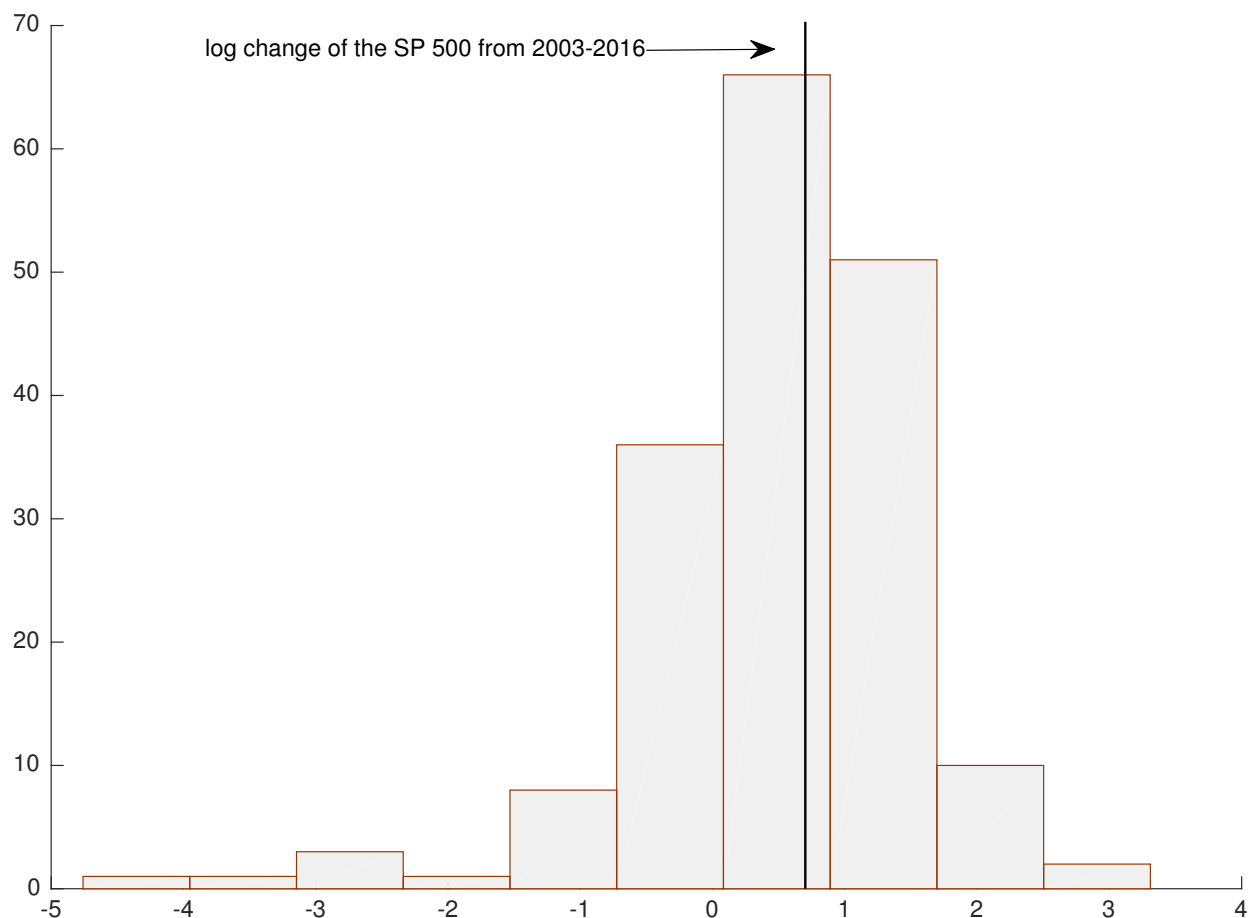


Figure 5: logarithm of the (gross) wealth growth rates of billionaire families between 2003 and 2016. Source: Forbes 2003 and Wealth-X 2016.

conclusion from the table is that, if anything, self-made individuals spend a slightly higher fraction of their wealth on philanthropy compared to billionaires who have inherited their wealth. Even though this observation pertains only to philanthropy, it is suggestive that differences in expenditure rates do not seem a likely candidate for the observed differences in wealth growth rates.

We summarize the pieces of the empirical evidence that motivates our model as follows: 1) Recent decades have seen an increase in portfolio allocations to alternative asset classes, real assets, and a simultaneous drop in the real interest rate. 2) The addition of new firms to the market portfolio acts in a manner similar to dilution for existing investors and over long periods is the dominant source of discrepancy between aggregate consumption and dividends per share of the market portfolio. 3) The rate of these additions has progressively increased in recent decades. 4) This increase coincided with a strong ascension to wealth for self-made




Figure 6: Kernel-smoothed density of log net worth for self made billionaires and billionaires who inherited their wealth. Sources: Forbes 2016 and Wealth-X 2016

	(1)	(2)	(3)	(4)	(5)	(6)
	lnphil	lnphil	lnphil	lnphil	lnphil	lnphil
lnw	1.157*** (0.151)	1.149*** (0.149)	1.199*** (0.341)	1.191*** (0.334)	0.983*** (0.166)	0.982*** (0.164)
SelfMade		0.234 (0.269)		0.0953 (0.517)		0.278 (0.313)
Constant	-8.312* (3.338)	-8.292* (3.313)	-9.691 (7.424)	-9.582 (7.324)	-4.160 (3.679)	-4.320 (3.640)
Observations	258	258	86	86	154	154
R^2	0.201	0.204	0.184	0.184	0.174	0.178

Table 2: Regressions of log(Philanthropy) on log(net worth) and a dummy variable taking the value one if the billionaire is listed as self made. Regressions (1) and (2) pertain to the entire sample, (3) and (4) to billionaires between 45 and 65 and (5) and (6) to billionaires above 65.

billionaires; by contrast the wealth growth of the 2003 billionaires (who are presumably more sophisticated than the typical investor) have largely followed the change in the value of the S&P, which has been far slower than the growth in the total wealth held by billionaires.

Based on these stylized facts, we build a model that departs from the standard assumption that the gains of innovation accrue to the representative agent. Instead we take seriously the notion that these gains are not equally shared between investors, and investigate the implications of such a departure for finance and macroeconomics.

3 Model

3.1 Agent preferences and demographics

We consider a model with discrete and infinite time: $t = \{\dots, 0, 1, 2, \dots\}$. The size of the population is normalized to one. At each date a mass λ of agents are born, and a mass λ dies so that the population remains constant. We denote by $V_{t,s}$ the utility at time t of an agent born at time s . Preferences are logarithmic:

$$\log V_{t,s} = \log c_{t,s} + \beta (1 - \lambda) \mathbf{E}_t \log V_{t+1,s} =: \log c_{t,s} + \hat{\beta} \mathbf{E}_t \log V_{t+1,s}, \quad (2)$$

where $\beta \in (0, 1)$ is the agent's subjective discount factor, and $c_{t,s}$ is the agent's consumption at time t . We also defined $\hat{\beta} = \beta(1 - \lambda)$ as the agent's effective discount factor. These preferences imply that the representative agent has an intertemporal elasticity of substitution (IES) equal to one and a risk aversion equal to one. These preferences are convenient for obtaining closed-form solutions. We consider extensions to allow for general risk aversion and IES later.

3.2 Technology

There is a continuum of intermediate-good firms that own non-perishable blueprints. These firms engage in monopolistic competition. We omit the details of the derivation, which is standard and can be found in Gârleanu et al. (2012), and state here the relevant elements for our asset-pricing purpose.

First, total output equals

$$Y_t = T_t A_t^{1-\alpha}, \quad (3)$$

where A_t is the number of intermediate goods available at time t and T_t is total factor productivity.

Second, an intermediate-good blueprint generates time- t profits

$$\pi_t = \frac{\alpha(1-\alpha)Y_t}{A_t} \propto A_t^{-\alpha}, \quad (4)$$

which implies

$$A_t\pi_t = \alpha(1-\alpha)Y_t. \quad (5)$$

In words, the key features of the model are that increasing the number of blueprints A_t raises total output Y_t (equation (3)), but the profits per blueprint decline (equation (4)). This is the sense in which this simple production specification captures the idea of displacement of old blueprints by new ones.

We would like to point out here that, even though we opted for a basic Romer-style production specification, the specific production assumptions (whether they are of the Romer type or the quality-ladder type) are irrelevant for the intuitions we develop in this paper.

3.3 New agents and products

The measure λ of newly born agents are of two types: a fraction θ are entrepreneurs and a fraction $1 - \theta$ are workers. Workers supply one unit of labor inelastically throughout their life. Since workers are not the focus of the paper, we assume that they are “hand-to-mouth” consumers, i.e., their wage income equals their consumption period-by-period. This assumption is not essential for the results, and we relax it in a later section.

Each period a total mass

$$\Delta A_{t+1} = A_{t+1} - A_t = \eta A_t (\Gamma_{t+1} + \Gamma_{t+1}^E) \quad (6)$$

of new blueprints arrives. Here, Γ_{t+1} is a gamma distributed variable with shape parameter a and rate parameter b , Γ_{t+1}^E is an independent gamma distributed variable with shape parameter a^E and rate parameter b^E , and η is a constant. All new blueprints arrive exogenously. The blueprints accounted for by Γ_{t+1}^E are distributed to existing firms.³ The other blueprints, accounted for by Γ_{t+1} , are assigned randomly to the measure $\lambda\theta$ of newly born entrepreneurs, as we describe immediately below.

One should view these new blueprints as ideas for the production of new products. The assumption that no resources are needed for the production of these blueprints or the implementation of these ideas (i.e., growth is exogenous) is for simplicity only, and we discuss

³Given our assumptions of frictionless trading of these firms, the allocations of blueprints is irrelevant.

how to relax it in a later section. The crucial aspect for our analysis is that the rents from the arrival of a subset of new blueprints are allocated heterogeneously to the newly arriving entrepreneurs, as follows.

At the time of their birth, a mass $\lambda\theta di$ of entrepreneurs is assigned to every “location” $i \in [0, 1]$ on a circle. Since the gamma distribution is infinitely divisible, we write the total number of blueprints distributed to entrepreneurs as $\eta A_t \Gamma_{t+1} = \eta A_t \int_{i \in [0,1]} d\Gamma_{i,t+1}$, where $d\Gamma_{i,t+1}$ denotes the independent increments of a gamma process and captures the increment in the mass of blueprints that will be assigned to location i on the circle at time $t + 1$.⁴

Since the gamma process is not commonly used in economics, we summarize briefly some of its properties. To build intuition, we consider a discrete construction. We split the interval $[0, 1]$ into N equal intervals, and think of the gamma process at the location $\frac{k}{N}$ as a sum of gamma-distributed increments $\xi_{\frac{n}{N}}$,

$$\sum_{n=1}^k \xi_{\frac{n}{N}}, \quad (7)$$

where the pdf of the increment ξ_i is given by

$$Pr(\xi_i \in dx) = \frac{b^{\frac{a}{N}}}{\Gamma\left(\frac{a}{N}\right)} x^{\frac{a}{N}-1} e^{-bx} dx. \quad (8)$$

The parameters $\frac{a}{N}$ and b are sometimes referred to as the “shape” and the “rate” of the gamma distribution, and $\Gamma\left(\frac{a}{N}\right)$ is the gamma function evaluated at $\frac{a}{N}$. The increments ξ_i are independent of each other, and the properties of the gamma distribution imply that

$$Pr\left(\sum_{n=1}^N \xi_{\frac{n}{N}} \in dx\right) = \frac{b^a}{\Gamma(a)} x^{a-1} e^{-bx} dx, \quad (9)$$

which is the distribution of a gamma variable with shape a and rate b .

Using the gamma process is technically attractive for our purposes, since it captures in a stylized way the fact that entrepreneurship is very risky. This is illustrated in Figure 7. The figure shows a sample of increments ξ_i for the case $N = 1000$. The figure illustrates that these increments tend to be close to zero for most of the locations; however, a small subset of random locations exhibit big spikes of random height. From an economic point of view, this means that only the lucky few entrepreneurs who happen to find themselves in the locations exhibiting the large spikes obtain valuable allocations of blueprints.

⁴For technical reasons, we think of new entrepreneurs as indexed by $(i, j) \in [0, 1] \times [1, 1]$, with, for all j , (i, j) assigned to location i and receiving the same number of blueprints $\eta A_t d\Gamma_{i,t+1}$.




Figure 7: An illustration of the increments ξ_i , for the case $N = 1000$, $a = 1$, and $b = 2$.

The limit of the variables given by (7) as the number of locations N goes to infinity is a gamma process. It is a positive and increasing process, whose paths are not continuous (they are only right continuous with left limits).⁵ This implies that any given location is likely to receive a negligible allocation of blueprints. However, because the process $\Gamma_{i,t+1}$ is a discontinuous function of i , a zero-measure of locations receive a strictly positive measure of blueprints, and the entrepreneurs who find themselves in these locations become spectacularly wealthy.

Before proceeding, we would like to note that this extreme-inequality setup is mostly for illustrative purposes and technical convenience. Less extreme distributions⁶ would not affect the economic insights, as long as we preserve some notion of distributional risk.

A final crucial assumption is that no agent knows at time t the realization of the path of the gamma process Γ_{t+1} . Everyone is trading behind the “veil of ignorance” regarding which locations on the circle will obtain the valuable blueprints and which ones will obtain the useless ones. Newly arriving entrepreneurs are therefore eager to share that risk by selling shares to investors on the market before this uncertainty is resolved. These shares entitle

⁵The gamma process is, however, continuous in probability. This means that, for any $\varepsilon > 0$, $\lim_{\delta \rightarrow 0} \Pr(|\Gamma_{i+\delta,t+1} - \Gamma_{i,t+1}| > \varepsilon) = 0$.

⁶Something as simple as assuming that locations are finite rather than a continuum would produce less extreme distributions without affecting the economic intuitions.

investors to a fraction v of the profits that will be produced by the newly arriving firms in perpetuity. A fraction $1 - v$ is “inalienable,” a reduced form way of capturing incentive effects of equity retention.

3.4 Markets

At each point in time, an investor can trade a zero net-supply bond. We follow Blanchard (1985) and assume that agents can also trade annuities with competitive insurance companies that break even. These annuity contracts entitle an insurance company to collect the wealth W_t^j of an agent j in the event that she dies at time t and in exchange provide her with an income stream λW_t^j while she is alive. We refer to Blanchard (1985) for further details.

Investors at time t can trade costlessly in the shares of all companies created prior to time t . Per blueprint, all such companies make the same profits π . For future use, we denote by Π_t the value of the future stream of profits from the representative blueprint:

$$\Pi_t = E_t \left[\sum_{t+1}^{\infty} \frac{M_s^i}{M_t^i} \pi_s \right], \quad (10)$$

with M_s^i the marginal-utility process of a given investor.

The shares of new firms are introduced to the market via intermediaries as follows: Every investor at time t is assigned to a location i on the circle $[0, 1)$. At each location i there is a representative competitive intermediary, who purchases (from the entrepreneurs) an equally weighted portfolio of the shares of the incipient companies located in an arc of length Δ_i centered at location i . The intermediary then offers the portfolio for purchase to the investors in location i . Intermediation requires resources equal to $\psi Y_t f(\Delta_i)$ per share. (We make the cost proportional to Y_t to ensure that the cost of intermediation is a stationary fraction of the size of the aggregate economy.) Hence, to break even, the intermediary needs to sell each share of the portfolio at a price $\frac{1}{\Delta_i} \int_{i-\frac{\Delta_i}{2}}^{i+\frac{\Delta_i}{2}} P_t^j dj + \psi Y_t f(\Delta_i)$, where P_t^j is the price of a newly created firm in location j . Assuming that there exist location-invariant equilibria such that $P_t^j = P_t$, the price of a portfolio share is simply $P_t + \psi Y_t f(\Delta_i)$. To simplify notation, from now on, we guess that there exist equilibria with $P_t^j = P_t$, and will then verify their existence in the next section. Likewise, we conjecture that Δ_i is independent of i , and write $\Delta_i = \Delta$. For further convenience, we let $\hat{P} = P_t + \psi Y_t f(\Delta)$.

Figures 8 and 9 illustrate how intermediaries can facilitate risk sharing in this economy. By purchasing an equal-weighted portfolio of shares on an arc of length Δ , the intermediaries are able to “smooth out” the spikes of the gamma process. Indeed, as the figure




Figure 8: An illustration of how intermediaries help with risk sharing. The increments are the same as in Figure 7, and $\Delta = 0.5$. The intermediary in position 0.5 provides an equal weighted portfolio of the increments in $[0.25, 0.75]$. The intermediary in position 0.75 averages the increments in $[0.5, 1]$, while the intermediary in position 0 averages the increments in $[0, 0.25] \cup [0.75, 1]$.

illustrates, they can offer their investors a portfolio of blueprints that has the same mean as the number of blueprints that arrive in each location, but is second-order stochastically dominant. Specifically, by using properties of the gamma distribution, one can show that $\frac{1}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} d\Gamma_j$ is gamma distributed with shape $a\Delta$ and rate $b\Delta$, and accordingly it has mean equal to $\frac{a}{b}$ and standard deviation equal to $\frac{\sqrt{a}}{b\sqrt{\Delta}}$. Further, it holds that, if $\Delta_2 > \Delta_1$, then $\frac{1}{\Delta_2} \int_{i-\frac{\Delta_2}{2}}^{i+\frac{\Delta_2}{2}} d\Gamma_j \succ^2 \frac{1}{\Delta_1} \int_{i-\frac{\Delta_1}{2}}^{i+\frac{\Delta_1}{2}} d\Gamma_j$, where \succ^2 denotes second-order stochastic dominance.

The arc-length Δ implies a specific correlation between the blueprints obtained by any given fund and the total growth of blueprints between t and $t+1$. Specifically, the correlation between the blueprints accruing to a given fund $\frac{1}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} d\Gamma_j$ and total new blueprints $\int_0^1 d\Gamma_j$ is $\sqrt{\Delta}$. Given the location invariance of the set-up, this correlation does not depend on i .

Intermediaries in each location are competitive and in an effort to attract investors they determine Δ in a way that maximizes investor welfare. Moreover, the assumption of perfect competition ensures that intermediaries make no profits. Accordingly, they act as simple




Figure 9: The distribution of equal weighted returns. The left figure depicts the blueprints accruing to the portfolio formed by the intermediary in each location i , which is simply an equal weighted average of the blueprints accruing to locations in an arc Δ around the intermediary’s location. The right figure is identical to the left figure except that the results are now depicted in polar coordinates.

pass throughs, enabling existing investors access to the newly created firms albeit at a cost.

To allow for heterogeneity in the returns of existing investors, we assume that $f(1) = \infty$, implying that Δ lies in the interior of $[0, 1]$.

We conclude this section with two comments. First, the notion of a “location” should not be understood geographically. It is simply a convenient device to produce heterogenous returns across investors. Second, since the gamma process has independent increments, it is immaterial whether investors invest in a single arc of length Δ or a set of non-contiguous arcs of total length Δ . Our construction only requires that these sets satisfy rotational symmetry.

3.5 Budget constraints

With these assumptions, the dynamic budget constraint of an investor who resides in location i can be expressed as

$$W_t^i = S_t^{E,i} A_t \Pi_t + B_t^i + S_t^{N,i} \hat{P}_t + c_t^i, \tag{11}$$

$$W_{t+1}^i = S_t^{E,i} A_t (\Pi_{t+1} + \pi_{t+1}) + (1 + r_t^f) B_t^i + S_t^{N,i} \frac{\eta v}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} (\Pi_{t+1} + \pi_{t+1}) A_t d\Gamma_{i,t+1} + \lambda W_{t+1}^i,$$

where $S_t^{E,i}$ is the number of shares of the (representative) firm that is already traded at time t , B_t^i is the amount invested in bonds, r_t^f the interest rate, and $S_t^{N,i}$ is the number of shares purchased in the intermediary-provided portfolio of newly created firms. We note that W_{t+1}^i is the agent's wealth at $t + 1$ conditional on survival. We normalize the supply of shares of all firms to unity. A convenient way to express (11) is

$$\frac{W_{t+1}^i}{W_t^i} = \left(\frac{1 - \frac{c_t^i}{W_t^i}}{1 - \lambda} \right) \left(\phi_B^i (1 + r_t^f) + \phi_E^i R_{t+1}^E + \phi_N^i R_{t+1}^{N,i} \right), \quad (12)$$

where $\phi_B^i \equiv \frac{B_t^i}{W_t^i - c_t^i}$, $\phi_E^i \equiv \frac{S_t^{E,i} A_t \Pi_t}{W_t^i - c_t^i}$, and $\phi_N^i = \frac{S_t^{N,i} \hat{P}_t}{W_t^i - c_t^i}$ are the post-consumption wealth shares invested by investor i in bonds, existing firms, and newly arriving firms respectively, and $R_{t+1}^E \equiv \frac{\Pi_{t+1} + \pi_{t+1}}{\Pi_t}$ and $R_{t+1}^{N,i} \equiv \frac{A_t \eta \nu}{\hat{P}_t \Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} (\Pi_{t+1} + \pi_{t+1}) d\Gamma_{i,t+1}$ are the gross returns of existing firms and the portfolio of newly arriving firms that investor i invests in.

An important observation about (12) is that, as long as $P_t^j = P_t$ for all $j \in [0, 1)$, the portfolio choices ϕ_B^i , ϕ_E^i , and ϕ_N^i , as well as the choices of Δ_i and of $\frac{c_t^i}{W_t^i}$ are the same for all investors, irrespective of their level of wealth and the location where they reside at time t , which simplifies the solution and analysis of the model.

To ensure that $P_t^j = P_t$ for all $j \in [0, 1)$ we make one final assumption, namely that investors re-locate prior to the start of trading each period, so that the total wealth of all the investors positioned in each location becomes equal across locations.⁷

This simplifying assumption is consistent with free movement of investors among locations. Conditional on all funds offering the same arc-length $\Delta_i = \Delta$, and given the location-invariant nature of the distribution of new firms across the circle, the investors have the incentive to position themselves in locations that offer lower prices for a share to the portfolio of new firms, thus equalizing these prices across locations. This outcome occurs when wealth moves across locations in such a way that the total wealth in every location is equalized.

3.6 Location-invariant equilibrium

The definition of a location-invariant equilibrium is standard. Such an equilibrium is a collection of prices Π_t , p_t , and P_t , portfolio allocations ϕ_B , ϕ_E , and ϕ_N , a choice of Δ , and consumption processes for all agents c_t^j such that a) Given prices, ϕ_B , ϕ_E , ϕ_N , Δ , and c_t^j are choices that maximize (2) subject to (12), b) the consumption market clears:

⁷Mathematically, such a re-location is always possible; one of the infinitely many ways to achieve it is to assign the investor with wealth W_t^j to location $F^{-1}(W_t^j)$, where $F(\cdot)$ is the wealth distribution.

$\int_j dc_t^j = A_t \pi_t - \psi T_t A_t^{1-\alpha} f(\Delta)$, c) the markets for all shares (both new and existing) clear: $\int_j dS_t^{E,j} = \int_j dS_t^{N,j} = 1$, and d) the bond market clears: $\int_j dB_t^j = 0$.

4 Solution

Next we construct an equilibrium that is both location-invariant, time-invariant and symmetric, in the sense that all agents choose the same portfolio. Specifically, we conjecture that there exists an equilibrium whereby $\phi_B = 0$, and the portfolio shares ϕ_E and ϕ_N , the interest rate r^f , the participation arc Δ , the valuation ratios $P^E \equiv \frac{\Pi_t}{\pi_t}$, $P^N \equiv \frac{P_t}{A_t \pi_t}$, and $\hat{P}^N \equiv \frac{\hat{P}_t}{A_t \pi_t}$, and the consumption-to-wealth ratio $c \equiv \frac{c_t^i}{W_t^i}$ are the same for all agents and constant across time. After computing explicit values for the constants that support such an equilibrium, we provide sufficient conditions for its existence.

For the remainder of this section we specialize the model to the case of logarithmic preferences. Later we show how to extend the results to the case of recursive preferences with general risk aversion, a case that we also use as basis for our quantitative evaluation.

Maintaining the supposition that the price-to-earnings ratio for existing firms $\frac{\Pi_t}{\pi_t}$ is a constant, denoted by P^E , the return R_{t+1}^E can be expressed as

$$\begin{aligned} R_{t+1}^E &= \frac{\pi_{t+1} + \Pi_{t+1}}{\Pi_t} (1 + \Gamma_{t+1}^E) = \frac{\pi_{t+1}}{\pi_t} \left(\frac{1 + P^E}{P^E} \right) (1 + \Gamma_{t+1}^E) \\ &= \frac{T_{t+1}}{T_t} \left(\frac{1 + P^E}{P^E} \right) \left(\frac{A_{t+1}}{A_t} \right)^{-\alpha} (1 + \Gamma_{t+1}^E) \end{aligned} \quad (13)$$

$$= \frac{T_{t+1}}{T_t} \left(\frac{1 + P^E}{P^E} \right) (1 + \eta (\Gamma_{t+1}^E + \Gamma_{t+1}))^{-\alpha} (1 + \Gamma_{t+1}^E). \quad (14)$$

Moreover, using similar reasoning, the return on an equal-weighted portfolio of newly created firms over an arc Δ is

$$R_{t+1}^N = \frac{v A_t (\Pi_{t+1} + \pi_{t+1}) \frac{\eta}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} d\Gamma_{i,t+1}}{A_t \pi_t \hat{P}^N} \quad (15)$$

$$= \frac{R_{t+1}^E}{1 + \Gamma_{t+1}^E} \frac{P^E}{\hat{P}^N} \frac{\eta v}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} d\Gamma_{i,t+1} \quad (16)$$

$$=: \frac{P^E}{\hat{P}^N} R_{t+1}^E Z_{t+1}. \quad (17)$$

where we defined the variable Z_{t+1} that captures the wedge between the two returns:

$$Z_t := (1 + \Gamma_{t+1}^E)^{-1} \frac{\eta v}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} d\Gamma_{i,t}. \quad (18)$$

Note that we omit the dependence of the random variables R_{t+1}^N and Z_{t+1} on i , since all the statements we make pertain to the joint distribution of this variable and aggregate quantities, which is independent of i .

The following proposition contains an explicit description of a symmetric time- and location-invariant equilibrium.

Proposition 1 *Assuming that a location-invariant, time-invariant, and symmetric equilibrium exists, the unique equilibrium values of ϕ_B , ϕ_E , ϕ_N , and c are*

$$\phi_B = 0 \tag{19}$$

$$\phi_E = E[(1 + Z_t)^{-1}] \tag{20}$$

$$\phi_N = 1 - \phi_E \tag{21}$$

$$c = 1 - \hat{\beta}. \tag{22}$$

With $\delta = \frac{\psi Y_t}{A_t \pi_t}$, the equilibrium values of P^E and \hat{P}^N are given by

$$P^E = \phi_E (1 - \delta f(\Delta)) \frac{\hat{\beta}}{1 - \hat{\beta}}, \tag{23}$$

$$\hat{P}^N = (1 - \phi_E) (1 - \delta f(\Delta)) \frac{\hat{\beta}}{1 - \hat{\beta}}, \tag{24}$$

and the interest rate equals

$$1 + r^f = \frac{E[(1 + Z_t)^{-1}]}{E[(R_t^E)^{-1} (1 + Z_t)^{-1}]}. \tag{25}$$

Finally, the equilibrium value of Δ is given by the solution to the equation

$$\frac{1 - \hat{\beta}}{\hat{\beta}} \frac{\delta f'(\Delta)}{1 - \delta f(\Delta)} = \frac{\partial E[\log(1 + Z_t)]}{\partial \Delta}. \tag{26}$$

We analyze the properties of the equilibrium in steps. First, we derive the implications of the equilibrium for risk sharing both within and across cohorts of entrepreneurs. Then we discuss implications for the size of the financial industry and its relation to the equilibrium expected excess returns of existing firms and new ventures.

4.1 Risk-sharing implications

For presentation purposes, it is convenient to start the analysis by treating Δ not as a choice variable, but rather as an exogenous parameter.

To derive the implications of the model for risk sharing between and across investor cohorts, we start with the following lemma.

Lemma 1 *Aggregate wealth growth is given by*

$$\frac{W_{t+1}}{W_t} = \frac{T_{t+1}}{T_t} (1 + \eta\Gamma_{t+1} + \eta\Gamma_t^E)^{1-\alpha}, \quad (27)$$

while an individual investor's wealth growth (conditional on survival) is given by

$$\frac{W_{t+1}^i}{W_t^i} = \frac{W_{t+1}}{W_t} \left(\frac{1}{1-\lambda} \right) \left(\frac{1+P^E}{1+P^E+P^N} \right) \left(\frac{1+\eta\Gamma_{t+1}^E + \eta v\Gamma_{t+1}}{1+\eta\Gamma_{t+1}^E + \eta\Gamma_{t+1}} \right) X_{i,t+1}, \quad (28)$$

where

$$X_{i,t+1} \equiv \frac{1 + \eta\Gamma_{t+1}^E + \eta v\Gamma_{t+1} \frac{1}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} dL_{j,t+1}}{1 + \eta\Gamma_{t+1}^E + \eta v\Gamma_{t+1}} \quad (29)$$

$$dL_{j,t+1} \equiv \frac{d\Gamma_{j,t+1}}{\Gamma_{t+1}}. \quad (30)$$

Equation (28) reveals that risk is imperfectly shared both within and across cohorts. The lack of within-cohort risk sharing is captured by the term $X_{i,t+1}$, which reflects heterogenous investment returns experienced by existing agents. This is driven by their inability to invest in all available ventures. Indeed, for $\Delta = 1$ the term $X_{i,t+1}$ becomes one, and the within-cohort lack of risk sharing disappears.

However, this is not the only dimension along which risk is imperfectly shared. Even if $\Delta = 1$, equation (28) shows that individual wealth $\frac{W_{t+1}^i}{W_t^i}$ and aggregate wealth $\frac{W_{t+1}}{W_t}$ are not perfectly correlated as long as v is less than one. The random term $\frac{1+\eta\Gamma_{t+1}^E + \eta v\Gamma_{t+1}}{1+\eta\Gamma_{t+1}^E + \eta\Gamma_{t+1}}$ captures the inter-cohort lack of risk sharing. It is driven by the fact that newly arriving entrepreneurs retain a fraction $1 - v$ of new company shares. Indeed, with $v = 1$ the term $\frac{1+\eta\Gamma_{t+1}^E + \eta v\Gamma_{t+1}}{1+\eta\Gamma_{t+1}^E + \eta\Gamma_{t+1}}$ disappears.

In summary, Δ controls the extent of intra-, while v controls the extent of inter-cohort risk sharing. If $\Delta = v = 1$, then risk is perfectly shared both within and across cohorts; individual wealth growth and aggregate wealth growth are perfectly correlated. However, even in that case individual and aggregate wealth growth differ by a negative constant. Indeed, aggregating the wealth growth of all investors surviving into $t + 1$, we obtain

$$\log \left((1-\lambda) \frac{\int_i W_{t+1}^i di}{W_t} \right) - \log \left(\frac{W_{t+1}}{W_t} \right) = \log \left(\frac{1+P^E}{1+P^E+P^N} \right) < 0. \quad (31)$$

The negative constant reflects that the wealth owned by existing investors does not include new-firm endowments, which belong to new entrepreneurs rather than existing investors.

4.2 Implications for the SDF

Since the wealth-to-consumption ratio is constant, our conclusions on wealth changes apply without modification to consumption changes of individual investors: an individual investor's consumption change is given by the right hand side of (28). With logarithmic utilities, the SDF M_t of an individual investor is given by

$$\begin{aligned} \frac{M_{t+1}^i}{M_t^i} &= \beta (1 - \lambda) \left(\frac{W_{t+1}^i}{W_t^i} \right)^{-1} \\ &\propto (1 + \eta \Gamma_{t+1}^E + \eta \Gamma_{t+1})^\alpha \left(1 + \eta \Gamma_{t+1}^E + \frac{\eta \nu}{\Delta} \Gamma_{t+1} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} dL_{j,t+1} \right)^{-1} \left(\frac{T_{t+1}}{T_t} \right)^{-1}. \end{aligned} \quad (32)$$

For the markets where all investors are participating (in particular, the market for existing stocks and the risk-free asset), any $\frac{M_{t+1}^i}{M_t^i}$ is a valid SDF, and so is

$$\frac{M_{t+1}}{M_t} \equiv \mathbb{E} \left[\frac{M_{t+1}^i}{M_t^i} \middle| \Gamma_{t+1} \right]. \quad (33)$$

By the properties of gamma distributed variables, the quantity $\int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} dL_{j,t+1}$ is beta distributed and independent of Γ_{t+1} . One can provide an explicit closed form solution for $\frac{M_{t+1}}{M_t}$ as a function of Γ_{t+1} in terms of a hypergeometric function. For our purposes, the interesting property of $\frac{M_{t+1}}{M_t}$ is its covariance with the growth shock Γ_{t+1} .

Lemma 2 *When risk sharing both across and within cohorts is perfect, i.e., $\Delta = 1$ and $\nu = 1$, the SDF $\frac{M_{t+1}}{M_t}$ is decreasing in Γ_{t+1} , and therefore $\text{Cov}\left(\frac{M_{t+1}}{M_t}, R_{t+1}^E\right) > 0$, if T is deterministic. However, $\text{Cov}\left(\frac{M_{t+1}}{M_t}, \Gamma_{t+1}\right) > 0$ and $\text{Cov}\left(\frac{M_{t+1}}{M_t}, R_{t+1}^E\right) < 0$ if either ν or Δ is sufficiently small.*

Lemma 2 shows how risk sharing imperfections can determine whether the marginal utility of consumption of the representative investor (broadly) rises or declines as the displacement innovation Γ_{t+1} increases. If risk is shared perfectly both within and across cohorts, then large realizations of Γ_{t+1} are “good news” for the representative investor. The gains in the value of the portfolio of new firms are enough to undo the losses on the existing assets owned by the investor. However, away from the perfect risk-sharing limit, large realizations of Γ_{t+1} are “bad news.” For instance, if risk is shared perfectly within cohorts ($\Delta = 1$) but imperfectly across cohorts ($\nu < 1$), then a fraction of the value of new ventures cannot be separated from the newly arriving cohort of agents. Hence, the losses on the portfolio of

existing assets cannot be offset by the gains on the portfolio of new ventures. This is the key insight of Gârleanu et al. (2012).

Even if risk is perfectly shared across cohorts ($v = 1$), large realizations of Γ_{t+1} may be (unconditionally) perceived as states of high marginal utility (“bad states”) when Δ is sufficiently below one. In this situation existing investors *as a group* gain from increased innovation, since they buy all the shares of the newly arriving entrepreneurs before the realization of Γ_{t+1} . However, the investors do not know *ex ante* whether they will receive a large or a small allotment of the new firms. Because of their risk aversion, they assign greater weight to the event that they end up with a disproportionately small share of the gains from growth, and therefore they perceive a high realization of Γ_{t+1} as bad news. This intuition is reminiscent of the intuition put forth by Constantinides and Duffie (1996), Kogan et al. (2019), and Gârleanu et al. (2015).

Hence our model nests models of perfect risk sharing as well as of imperfect risk sharing, across and within cohorts, as special cases. However, the most important difference is that it proposes a view of the financial industry as a (costly) device to improve risk sharing, yielding joint predictions on how expected returns, the size of the financial industry, the interest rate, etc., change as, say, displacement increases.

4.3 Equilibrium excess returns

Having derived the equilibrium SDF, we can now discuss the implications of the model for expected excess returns. We start with a comparison of the expected returns on new ventures as opposed to existing firms. We then discuss implications for the excess return on the stock market. Using (14), (17), and the results of Proposition 1 leads to

$$E \left[\frac{R_t^N}{R_t^E} \right] = \frac{1 - \phi_N}{\phi_N} \times E[Z_t]. \quad (34)$$

The right-hand side of (34) contains two terms. The first term, $\frac{1 - \phi_N}{\phi_N}$, is a declining function of ϕ_N , the share of aggregate wealth invested in new ventures. The second term has expectation equal to $\eta v \frac{a}{b} (E[(1 + \eta \Gamma_t^E)^{-1}])$, an expression that features only exogenous variables. Equation (34) shows that the expected return on new investments as compared to existing investments is negatively related to ϕ_N .

Lemma 3 *For any $\Delta \in [0, 1]$ and any $v \in [0, 1]$ it holds that $E \left[\frac{R_t^N}{R_t^E} \right] \geq 1$.*

Lemma 3 may appear counterintuitive at first pass. When ϕ_N is small, the bulk of investors’ wealth is invested in existing assets, whose value is declining in aggregate displacement Γ_{t+1} . Accordingly, one would expect new assets to act as a partial hedge against

capital losses on existing assets, and hence to have a low expected return. To see this, rewrite (17) as

$$R_{i,t+1}^N = \frac{1 + P^E}{\hat{P}^N} \frac{v\eta\Gamma_{t+1}}{(1 + \eta\Gamma_{t+1}^E + \eta\Gamma_{t+1})^\alpha} \frac{1}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} dL_{j,t+1}.$$

The properties of the gamma process imply that $\frac{1}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} dL_{j,t+1}$ is independent of Γ_{t+1} , and thus $R_{i,t+1}^N$ can be viewed as containing the component $\frac{v\eta\Gamma_{t+1}}{(1+\eta\Gamma_{t+1}^E+\eta\Gamma_{t+1})^\alpha}$, which is increasing in Γ_{t+1} , and the “idiosyncratic” component $\frac{1}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} dL_{j,t+1}$. The investor values the hedging component of $R_{i,t+1}^N$; absent other effects, its equilibrium excess return would be low, not high.

The explanation of the lemma lies with the idiosyncratic risk that the investors must take. Taking the limit as Δ approaches zero is useful: In that limit, both the variance and the skewness of $\frac{1}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} dL_{j,t+1}$ become arbitrarily large, while the expected value is always one. In economic terms, this means that an investment in new assets results in a loss with probability approaching one, and in a spectacularly high return with very small probability, which makes it unattractive for a risk averse investor.

At the opposite extreme of perfect risk sharing ($\Delta = 1, v = 1$), the expected excess return of new ventures continue to exceed the expected excess of existing assets, but for a different reason: in this case the marginal agent’s consumption is perfectly correlated with aggregate consumption and new ventures deliver high payoffs (a large number of blueprints) in states of the world in which the consumption growth of existing investors would be high, not low. Differently put, the pricing kernel decreases in Γ_{t+1} .

In summary, for any constellation of parameters, the expected ratio of gross returns of new ventures to existing assets exceeds unity. Interestingly, this gap exists even in cases where the investor’s marginal utility increases in Γ_{t+1} , i.e., in cases where the investor desires to hedge against fluctuations in displacement risk. Indeed, the smaller is Δ — and hence the higher the desire to hedge displacement risk — the larger is the gap in expected returns between the new venture portfolio and existing firms.

Inspection of (14) and (15) reveals that the returns R_{t+1}^E and R_{t+1}^N are negatively correlated. This result hinges critically on the fact that — purely for simplicity — the model includes only displacement shocks, but no “neutral” productivity shocks (or any other business cycle shocks for that matter), which would affect both new and existing firms in a similar fashion. Such shocks would easily render the correlation between R_{t+1}^E and R_{t+1}^N positive, but less than one; moreover they would introduce an additional source of risk for both existing

and new assets, which would require additional compensation in form of an enhanced excess return.

4.4 Equilibrium interest rate

For small enough Δ and v , and provided that α is not too small — i.e., displacement is not trivial — the equilibrium interest rate is declining in η , and is increasing in v and Δ .

A higher value of η implies that new assets amount to a larger fraction of existing stock market valuation. It also implies a more rapid (and more volatile and skewed) decline in the value of existing assets. Faced with such a threat to the value of their existing asset portfolio, investors see an increased need to invest in the risk-free asset. Since the risk-free asset is in zero net supply, this results in a decline in the real interest rate. Interestingly, the interest rate declines even though expected aggregate growth goes up. Clearly, this outcome is due to the fact that the gains from aggregate growth are unequally shared across investors, to the point that some investors may make losses.

An increase in v implies that existing agents can use purchases of new assets to reduce the impact of displacement shocks on their existing assets. Hence, less precautionary savings are required, and the interest rate increases. Finally, an increase in Δ also helps to reduce the idiosyncratic volatility of an investor's wealth, and hence helps reduce precautionary savings.

We note here that an increase in pure redistribution risk, which could be modelled as a lengthening of the circumference of the circle, coupled with a reduction in η so as to keep the expected arrival of new blueprints fixed, is mathematically isomorphic to a reduction in Δ keeping the circumference of the circle at one. Hence, an increase in pure redistribution risk has the effect of reducing the interest rate.

4.5 The participation arc Δ and the size of the financial industry

So far we have treated the participation arc Δ as an exogenous parameter. Now we discuss how it is determined inside the model. Equation (26) determines the size of the participation arc Δ , which is a monotone function of the resources devoted to the financial industry.

To gain some intuition on the determinants of the size of the financial industry, we provide the following comparative-static results.

Lemma 4 *Δ is an increasing function of η and v and a declining function of δ .*

Lemma 4 states that, as expected displacement increases (an increase in η), it becomes more attractive to expend resources to reduce the uncertainty associated with risky new ventures. Similarly, the lower the fraction of shares that is retained by newly arriving agents (v), the larger the incentive to expend resources to risk-share with the newly arriving agents. Finally, a higher value of δ increases the cost of the financial industry and hence the resources expended on it.

The following lemma summarizes the impact of increased displacement on the economy, taking into account the impact of endogenizing Δ .

Lemma 5 *Increased displacement (a higher η) implies a larger fraction of resources devoted to the financial industry (a higher Δ), and a higher fraction of the aggregate portfolio directed towards new ventures (higher ϕ_N).*

An alternative experiment would involve increasing the variability of A_{t+1}/A_t , and therefore the dispersion induced by displacement, while keeping its mean constant. One way to formalize that would be to decrease the gamma-distribution parameters a and b by the same factor.

5 Alternative investments

We next sketch how to extend the model to allow for a discussion of private equity and so called “real” assets such as real estate.

5.1 Private equity

To introduce private equity, we assume that each period a fraction of the existing firms lose their eligibility to receive new blueprints. Newly arriving entrepreneurs have the ability to purchase those firms from existing agents for a value Π per blueprint, and restore their ability to receive an allocation of blueprints over the next period, at which point the firms are re-introduced into the public market.

The private equity firms can sell shares to the investors with whom they share the location. Mirroring the assumptions of the baseline model, these firms can only purchase companies located within a distance of $\frac{\Delta}{2}$ from the firm.

This version of the model would be equivalent to the baseline model. Hence, whether the blueprints arrive to newly created firms, or to existing firms who occasionally lose the ability to receive new blueprints, at which point they are taken private and then re-listed,

is irrelevant for this model. In either version of the model, the benefits from the arrival of new blueprints would be asymmetrically distributed both across and within investor cohorts leaving the economic intuitions we identified unchanged.

5.2 Real assets

So far we have studied a model where labor is the only factor of production. We now extend the model to introduce additional factors. We distinguish between two types of factors, namely those that are not tied to a specific blueprint but are useful for all productive purposes, and those that are specific to a given blueprint. An example of the first factor of production would be commercial real estate or commodities, while an example of the second factor of production would be specialized equipment required to manufacture a given intermediate good.

5.2.1. Land

The introduction of a factor such as land is straightforward. To be as explicit as possible that land is not tied to any intermediate good, we assume that land is useful only in the production of the final good.

Land is owned by existing agents and rented out to final-good producing firms, so that aggregate output is given by

$$Y_t = F_t^\zeta (L_t^F)^{1-\alpha-\zeta} \left(\int_0^{A_t} x_{j,t}^\alpha dj \right), \quad (35)$$

where $\zeta \in (0, 1 - \alpha)$ is the share of output that accrues to land. Total land is fixed and normalized to one. Given the Cobb-Douglas structure of (35), it follows that the rental rate of land is

$$r_t^F = \zeta Y_t. \quad (36)$$

Once again, we construct an equilibrium where the price-to-rent ratio $P^F = \frac{P_t^F}{r_t^F}$ is constant, so that the return on land is given by $R_{t+1}^F = \frac{r_{t+1}^F + P_{t+1}^F}{P_t^F} = \frac{Y_{t+1}}{Y_t} \frac{1 + P^F}{P^F}$. Repeating the arguments of Section 4.1, the wealth evolution of an individual investor, conditional on

survival, is

$$\frac{W_{t+1}^i}{W_t^i} = \frac{1}{1-\lambda} \frac{W_{t+1}}{W_t} \left(\frac{\zeta(1+P^F)}{\alpha(1+P^E+P^N)+\zeta(1+P^F)} + \frac{\alpha(1+P^E)}{\alpha(1+P^E+P^N)+\zeta(1+P^F)} \left(\frac{1+\eta v \Gamma_{t+1}}{1+\eta \Gamma_{t+1}} \right) X_{i,t+1} \right), \quad (37)$$

for some new constants P^F , P^N , and P^E . Comparing (37) with (28), the only difference is that the wealth growth of an individual investor now gains a fraction $\frac{\zeta(1+P^F)}{\alpha(1+P^E+P^N)+\zeta(1+P^F)}$ of aggregate wealth growth. The reason is intuitive: Since land captures a constant fraction of total output, it actually benefits from higher values of Γ_{t+1} , since those are associated with higher output growth.

For sufficiently small Δ , v , and ζ , it follows that the SDF is negatively related to Γ_{t+1} . Since the return R_{t+1}^E is declining in Γ_{t+1} , while R_{t+1}^F is increasing in Γ_{t+1} , it follows that $E(R_{t+1}^F) < E(R_{t+1}^E)$. Indeed, in addition $E(R_{t+1}^F) < 1+r^f$, a result that, however, depends critically on the absence of neutral productivity shocks in the model.

Clearly, an increase in η will render investments in land more attractive, as a hedge to increased displacement risk.

5.2.2. Specialized equipment

If a factor of production was closely tied with a specific blueprint, its behavior would be quite different from that of a factor of production that is not. To illustrate what would happen in such a case, we assume that, in addition to labor, the production of the new goods requires a location-specific capital k_i :

$$x_{i,t} = k_i^\nu l_{i,t}^{1-\nu},$$

where $l_{i,t}$ is the amount of labor used in the production of intermediate good i and k_i denotes an irreversible capital investment k_i that is specific to the location of a blueprint and its vintage. Capital for the arriving vintages is produced by converting consumption goods to capital goods with one unit of investment good requiring one unit of the consumption good.

Similar to the baseline model, the financial industry allows existing investors at time Δ to invest in capital goods in an arc of length Δ without knowing which locations will be receiving blueprints in the next period. This capital is then sold to the newly arriving firms in the respective locations once production commences at time $t+1$. We suspend the market for the sale of new firm shares to existing investors prior to the resolution of the uncertainty

about their productivity; shares of new firms are tradeable only after their productivity is known. This assumption is inessential, but it will allow us to make a simple analogy to the baseline model.

We will only sketch the solution of this version of the model, since the key intuitions are no different than in the baseline model. In this version of the model aggregate output evolves according to

$$\frac{Y_{t+1}}{Y_t} = \left(\frac{A_{t+1}}{A_t} \right)^{1-(1-\nu)\alpha}$$

in steady state. The owners of capital goods extract a fraction ν of the present value of profits of the firms produced in location i , so that the return from investing in the new capital goods is given by

$$R_{t+1}^N = \frac{\nu A_t (\Pi_{t+1} + \pi_{t+1}) \frac{\eta}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} d\Gamma_{i,t+1}}{A_t k_t (1 + \psi f(\Delta))}, \quad (38)$$

which parallels closely equation (17), except that existing investors obtain a fraction ν of total profits and the cost of their investment is given by $A_t k_t (1 + \psi f(\Delta))$, where $A_t k_t$ is the number of investment goods and $1 + \psi f(\Delta)$ is the cost per unit of capital good. Since the equilibrium of this model features a constant price-to-profits ratio in equilibrium, equation (38) can alternatively be expressed in q-theoretic fashion as

$$R_{t+1}^N = \frac{\nu \frac{\pi_{t+1}}{k_t} (1 + P^E) \frac{\eta}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} d\Gamma_{j,t+1}}{1 + \psi f(\Delta)}, \quad (39)$$

where $\nu \frac{\pi_{t+1}}{k_t}$ is the marginal product of capital at firm inception and $1 + \psi f(\Delta)$ is the marginal cost of converting one unit of consumption to capital. Comparing (39) with (17) shows that ν plays a similar role to v , in that it controls the fraction of new firm value that accrues to existing investors. However, unlike the baseline model, the reason why existing investors capture that fraction is that they are providing the capital goods required for production, rather than insuring the entrepreneur.

In this version of the model the equilibrium quantity that adjusts to clear markets is k_{t+1} , the quantity of capital, rather than the valuation ratio P^N . Hence, the intuition of the baseline model that pertains to the magnitude of P^N carry over to the determination of k_{t+1} . For example, a pure increase in redistribution risk or an increase in ψ (a decrease in the efficiency of the financial industry) reduces the investment in new capital goods in steady state. By contrast, the attractiveness (and hence the equilibrium price) of a factor like land, which is not tied to a specific factor of production, increases under such conditions.

We conclude by noting that even though we drew a stark distinction between factors of production that are tied to specific blueprints and those that are not, in reality the distinction is more nuanced, with many factors being convertible between different uses at some cost.

6 Further Discussion

We discuss here how to interpret in the context of our model various other economic quantities and institutions.

6.1 Labor income and pension funds

Labor income benefits from displacement in our model since the arriving firms compete for labor services, which are in fixed supply. Indeed, total wages increase at the same rate as aggregate output. Moreover, we have assumed that workers are hand-to-mouth consumers who are not participating actively in financial markets. These assumptions are purely for simplicity and can be easily relaxed.

Suppose for instance that workers are allowed to participate in financial markets. Moreover, suppose that the production of a unit of each intermediate good takes one unit of a labor composite good, which is a Cobb-Douglas aggregate of labor inputs provided by different cohorts of workers

$$l_t = \prod_{s=-\infty..t} (l_{t,s})^{a_{t,s}}, \quad (40)$$

where $l_{t,s}$ is the labor input of workers born at time s and the weights $a_{t,s}$ are given by $a_{t,s} = \frac{\Delta A_s}{A_t}$. This specification implies that even though the aggregate wage bill grows at the rate of aggregate output, the fraction of wages accruing to a given cohort of workers declines over time and indeed at the same rate as the profits of existing firms.

One motivation for such a specification is skill obsolescence: As the number of blueprints expands, the skills of a given cohort of workers becomes progressively less useful.

If one were to adopt equation (40), and allowed workers access to financial markets on the same terms as firm owners, then all our conclusions would carry through without modification: With such a specification, workers' human capital would exhibit a similar exposure to displacement to that of the value of existing firms, making workers eager to hedge displacement risk by investing in newly arriving companies.

a^E	4	β	0.97	v	0.8
a	0.8	$1 - \lambda$	0.98	Δ	0.5
b	1.5	γ	6	α	0.92
δ	0.01	$f(\Delta)$	1		

Table 3: Parameters used for the calibration

In particular, if one took the view that pension funds invest on behalf of workers in a way that maximizes their welfare, then an increase in displacement activity would explain the increased popularity amongst pension funds of investment vehicles offering (positive) exposure to displacing firms.

7 Calibration (Preliminary)

In this section we calibrate the model. Our goal is to obtain a sense of the magnitudes of the share of the aggregate portfolio held in the form of alternative investments, the expected returns, interest rates, etc., that are implied by this model.

For calibration purposes it is convenient to maintain the assumption that agents have unit intertemporal elasticity of substitution, but have risk aversion γ that may be different from one. This feature doesn't change the model in a substantive way. The appendix describes the straightforward modifications necessary.

We abstract from issues such as illiquidity, lock-up periods, etc., which are quite common in private equity investments. As a result we choose to define a period to be equal to five years in our calibration. By lengthening the horizon, the Euler equation applies between times when the investor may actually have a chance to make a portfolio choice in reality. Consistent with common practice, we report returns at an annualized frequency to facilitate comparison with the literature.

Table 3 contains the parameters that we choose. The parameters a^E , a , and b control the distributions for the arrival of blueprints to existing and new firms. We choose these parameters to match as closely as possible the mean and variance of five-year real growth rates of aggregate consumption and the dividends of the market portfolio. To obtain a better understanding of these parameters, we note that the parameter a^E is five times larger than a . This means that (on average) the existing firms obtain five blueprints for every blueprint accruing to a newly born firm over a five-year horizon. As a result, the ratio $\frac{a^E}{a}$ controls how quickly the market value of existing firms declines as a fraction of total market capitalization

	$E(R^E)$	$E(R^N)$	r^f	ϕ_N
Baseline	3.54	13.40	1.76	6.48
$\Delta = 0.7$	2.08	5.39	3.34	8.48
$v = 0.5$	3.80	5.63	1.47	5.85
$\beta = 0.99$	3.47	13.10	-0.50	6.48
$\delta = 0.02$	3.55	13.43	2.04	6.48

Table 4: Un-levered excess return on existing assets ($E[R^E]$), un-levered excess return of the new assets portfolio provided by the intermediary ($E[R^N]$), risk-free rate (r^f), and fraction of assets invested in new assets (ϕ_N). All entries in the tables are percentages. All returns and excess returns are annualized. Rows correspond to different calibrations changing one parameter while keeping all others unchanged.

over a five-year horizon. With these choices the market capitalization of existing firms is on average 87% of total market capitalization at the end of the five-year period, a number that is consistent with the data. The parameter b is mostly a scaling constant. With this choice of b we obtain an average (annualized) consumption growth rate of about 2 percent per year and a volatility of consumption equal to 1.5% (annualized). The volatility of dividend growth of the market portfolio is approximately 7% (annualized). The correlation between dividend growth of the market portfolio (and accordingly stock market returns in this model) and consumption growth is approximately 10%.

We choose $1 - \lambda = 0.98$ to capture a birth rate of about 2% in the population, and $\beta = 0.97$ to match a real risk free rate of about 1%. The risk-aversion parameter $\gamma = 6$ is sufficient to match the (un-levered) equity premium. The parameter $\alpha = 0.92$ is chosen to match the profit share of output in aggregate data (about 8%). The quantity $\delta f(\Delta)$ captures the added value of the financial industry as a share of aggregate profits. Since the financial industry for the purposes of this paper is the private-equity industry, we choose a very low number. This number is actually quite inconsequential for any of the quantities of interest.

Finally, instead of fully specifying a function $f(\Delta)$ we choose Δ directly, since this is the quantity that enters returns, interest rates and the share of the market portfolio that is invested in alternative investments. A choice of Δ equal to 0.5 implies that a typical investor's portfolio of new firms has a correlation of $\sqrt{0.5}$ with the cross-sectional average of the returns of all private equity funds. Finally, the parameter $v = 0.8$ means that the




Figure 10: Left plot: Model-implied distribution of 5-year differences in log aggregate consumption. Right plot: Model-implied distribution of 5-year differences in log wealth of a typical investor (right panel). The vertical line depicts the model-implied 5-year difference in the log price of the market index.

founder retains approximately 20% of a firm’s equity.

Table 4 reports results of the calibration. The results are for un-levered returns. To relate un-levered to levered returns (which is what we observe in the data) one has to multiply the excess returns reported in the table by 1.6. Taking that into account, it is evident that the model can easily produce large equity premiums and also low real interest rates. Moreover, the fraction of assets invested in new assets is roughly consistent with the data.

The reason for the model’s quantitative success is quite simple: the representative investor’s wealth (and, accordingly, consumption) growth is more volatile than aggregate consumption. Figure 10 provides an illustration. An investor’s (annualized) standard deviation of consumption growth is approximately 5.7% in this model, while the annualized standard deviation of aggregate consumption growth is about 1.1%. The right plot of Figure 10 is the model analog of Figure 5, except that in the data the wealth changes are over a 13-year rather than a 5-year period. Comparing the two figures, we see that both imply a significant dispersion in wealth growth rates. If anything, the model implies a smaller dispersion of changes than the data, since the annualized standard deviation of wealth changes inside the model is about 6%, while in the data it is around 15%. Hence, the model’s asset pricing implications are not driven by extreme assumptions on the dispersion of wealth growth rates.

8 Conclusion

We proposed a tractable framework whereby the gains of growth are asymmetrically distributed across investors. The main results can be summarized as follows: An acceleration in displacement activity (and more generally the dispersion of the gains from growth) will increase heterogeneity amongst investors. The benefits will accrue predominantly to arriving entrepreneurs and lucky investors who happen to invest in profitable incipient firms. Existing equities will experience outflows in favor of asset classes such as fixed income, private equity, venture capital, and real assets. Increased growth may co-exist with lower interest rates. Low discount rates may not result in more physical investment but may simply result in an expansion of the financial industry.

Interestingly, the expected ratio of gross returns in incipient vs. existing firms may exceed one, even though incipient firms offer hedging benefits to investors. The reason is that in a world of imperfect risk sharing, idiosyncratic risk is priced.

Incomplete Reference List

References

Blanchard, O. J. (1985). Debt, deficits, and finite horizons. *Journal of Political Economy* 93(2), 223–247.

Constantinides, G. and D. Duffie (1996). Asset pricing with heterogeneous consumers. *Journal of Political Economy* 104, 219–240.

Gârleanu, N., L. Kogan, and S. Panageas (2012). The demographics of innovation and asset returns. *Journal of Financial Economics* 105, 491–510.

Kogan, L., D. Papanikolaou, and N. Stoffman (2019). Left behind: Creative destruction, inequality, and the stock market. *Journal of Political Economy*, forthcoming.

A Calibrated Model

For the purposes of the calibration, we adopt recursive preferences with general risk-aversion parameter γ , while keeping the intertemporal elasticity of substitution equal to one. As is well known, an agent's consumption-to-wealth ratio is the same as in the log-utility case, while her marginal utility growth, given the iid investment opportunity set, is proportional to $(W_{t+1}^i/W_t^i)^{-\gamma}$, and therefore, given (12), to

$$\left(\phi_B^i (1 + r^f) + \phi_E^i R_{t+1}^E + \phi_N^i R_{t+1}^{N,i}\right)^{-\gamma} = \left(\frac{P^E}{P^E + \hat{P}^N} R_t^E (1 + Z_t)\right)^{-\gamma}, \quad (41)$$

where the right-hand side is due to (46).

Similar computations to the ones in the proof of Proposition 1 yield

$$\phi_E = \frac{\mathbb{E} \left[(R_t^E)^{1-\gamma} (1 + Z_t)^{-\gamma} \right]}{\mathbb{E} \left[(R_t^E)^{1-\gamma} (1 + Z_t)^{1-\gamma} \right]} \quad (42)$$

and

$$1 + r^f = \frac{\mathbb{E} \left[(R_t^E)^{1-\gamma} (1 + Z_t)^{-\gamma} \right]}{\mathbb{E} \left[(R_t^E)^{-\gamma} (1 + Z_t)^{-\gamma} \right]}. \quad (43)$$

Finally, the equilibrium value of Δ is given by the solution to the equation

$$(\gamma - 1) \frac{1 - \hat{\beta}}{\hat{\beta}} \frac{\delta f'(\Delta)}{1 - \delta f(\Delta)} = -\frac{\partial}{\partial \Delta} \log \left(\mathbb{E} \left[(R_t^E)^{1-\gamma} (1 + Z_t)^{1-\gamma} \right] \right). \quad (44)$$

B Proofs

Proof of Proposition 1. A first-order condition for portfolio choice for an investor with logarithmic preferences is

$$\mathbb{E} \left[\frac{R_t^E - R_t^N}{\phi_B (1 + r^f) + \phi_E R_t^E + \phi_N R_t^N} \right] = 0. \quad (45)$$

Using the definitions of ϕ_E and ϕ_N and imposing $\phi_B = 0$ and market clearing in the stock markets implies $\phi_E = \frac{P^E}{P^E + \hat{P}^N}$ and $\phi_N = 1 - \phi_E$. Accordingly, using (14) and (17),

$$\phi_B (1 + r^f) + \phi_E R_t^E + \phi_N R_t^N = \frac{P^E}{P^E + \hat{P}^N} R_t^E (1 + Z_t). \quad (46)$$

Using (46) and (17) inside (45) and noting that $\frac{P^E}{\hat{P}^N} = \frac{\phi_E}{1 - \phi_E}$ leads to (20).

Having determined ϕ_E , it is straightforward to determine P^E and P^N . To start, we note that with logarithmic preferences $c = 1 - \beta(1 - \lambda)$. Aggregating (11) across agents and imposing asset market clearing implies

$$W_t = A_t \pi_t (P^E + P^N + 1). \quad (47)$$

Given goods market clearing $\frac{C_t}{A_t \pi_t} = (1 - \delta f(\Delta))$, we obtain

$$1 + P^E + P^N = \frac{1 - \delta f(\Delta)}{1 - \beta(1 - \lambda)}. \quad (48)$$

Combining (48) with $\phi_E = \frac{P^E}{P^E + \hat{P}^N}$ results in (23)–(24).

The first-order condition for the excess return $R_t^E - (1 + r^f)$ yields (25).

Finally, equation (26) follows from maximizing log utility over the right-hand side of (46), making use of the envelope theorem and recalling $\hat{P}^N = P^N + \delta f(\Delta)$. ■

Proof of Lemma 1. Equation (27) follows from (47) and (6). To arrive at equation (28), we use (12), (46), (14), and the goods market clearing condition $c \frac{W_t}{A_t \pi_t} + \delta f(\Delta) = 1$, which implies

$$1 - c = \frac{P^E + P^N + \delta f(\Delta)}{(1 + P^E + P^N)}.$$

■

Proof of Lemma 2. When $\nu = 1$ and $\Delta = 1$, $\frac{M_{t+1}}{M_t} \propto (1 + \eta \Gamma_{t+1}^E + \eta \Gamma_{t+1})^{\alpha-1}$, which is decreasing in Γ_{t+1} .

To prove the second sentence, note first that, for $\nu = 0$,

$$\frac{M_{t+1}}{M_t} \propto (R_{t+1}^E)^{-1} \propto (1 + \eta \Gamma_{t+1}^E)^{-1} (1 + \eta \Gamma_{t+1}^E + \eta \Gamma_{t+1})^\alpha \frac{T_t}{T_{t+1}}, \quad (49)$$

so that the desired statements hold obviously. By continuity in ν , they also hold for $\nu > 0$ small enough.

Note now that, for any $\nu > 0$ and $\Delta > 0$,

$$\frac{M_{t+1}^i}{M_t^i} \leq (1 + \eta \Gamma_{t+1}^E + \eta \Gamma_{t+1})^\alpha (1 + \eta \Gamma_{t+1}^E)^{-1} \left(\frac{T_{t+1}}{T_t} \right)^{-1}. \quad (50)$$

Furthermore, as $\Delta \rightarrow 0$, the probability that (50) holds with equality tends to zero. This probability is independent of Γ_{t+1}^E and Γ_{t+1} . It follows that, in the limit, the covariances with

M_{t+1}/M_t equal the respective covariances with the right-hand side of (50). The statement therefore holds for Δ small enough. ■

Proof of Lemma 3. $E \left[\frac{R^N}{R^E} \right]$ can be written as

$$E \left[\frac{R^N}{R^E} \right] = \frac{E \left[(1 + Z_t)^{-1} \right] E[Z_t]}{E \left[(1 + Z_t)^{-1} Z_t \right]} = 1 - \frac{\text{cov} \left((1 + Z_t)^{-1}, Z_t \right)}{E \left[(1 + Z_t)^{-1} Z_t \right]} > 1,$$

where the first equality follows from (20) and (34) and the inequality from the fact that Z_t and $(1 + Z_t)^{-1}$ are negatively correlated. ■

Proof of Lemma 4. It suffices to show that

$$\frac{\partial^2 E \left[\log \left(1 + \eta \Gamma_t^E + \frac{\eta v}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} d\Gamma_{i,t} \right) \right]}{\partial \Delta \partial \eta} > 0. \quad (51)$$

Differentiating first with respect to η gives

$$\frac{\partial E \left[\log \left(1 + \frac{\eta v}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} d\Gamma_{i,t} \right) \right]}{\partial \eta} = E \left[\frac{\Gamma_t^E + \frac{v}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} d\Gamma_{i,t}}{1 + \eta \Gamma_t^E + \frac{\eta v}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} d\Gamma_{i,t}} \right].$$

The function $x \mapsto \frac{x}{1+\eta x}$ is concave. Further, since Γ_t^E and $\frac{v}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} d\Gamma_{i,t}$ are independent and $\frac{v}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} d\Gamma_{i,t}$ increases in Δ in the sense of second-order stochastic dominance, so does $\Gamma_t^E + \frac{v}{\Delta} \int_{i-\frac{\Delta}{2}}^{i+\frac{\Delta}{2}} d\Gamma_{i,t}$. Inequality (51) follows. ■