Drifting apart: The pricing of assets when the benefits of growth are not shared equally*

Nicolae Gârleanu
UC Berkeley-Haas, NBER, and CEPR

Stavros Panageas
University of Chicago, Booth School of Business and NBER

Dimitris Papanikolaou
Northwestern University, Kellogg School of Management and NBER

Jianfeng Yu
University of Minnesota, Carlson School of Business

August 2015

Abstract

A significant fraction of the growth of aggregate market capitalization is due to new firm entry. With incomplete markets, the gains from new firm creation are not shared equally. Rather, they accrue to a small part of the population, and by potentially displacing existing firms constitute a risk for the marginal investor. We capture these notions in a simple model, and develop a methodology to measure the displacement risk, relying on the discrepancy in the growth rates of aggregate dividends and of the gains from the self-financing trading strategy associated with maintaining a market-weighted portfolio. We find that our measure of displacement risk is closely linked to certain cross-sectional asset-pricing phenomena and can explain a sizable fraction of the equity premium. We argue more generally that dispersion in capital income, a source of risk overlooked in representative agent models, has first-order implications for asset pricing.

Keywords: Asset Pricing, Displacement Risk, Equity Premium, Incomplete Markets, Innovation

*We are grateful to participants at the University of Chicago Finance lunch for comments. Jason Kang provided truly exceptional research assistance.

1 Introduction

New and young companies are an important source of job creation and economic growth in the American economy (Haltiwanger, 2012). However, the economic rents resulting from new firm entry are not distributed equally in the population, as evidenced by the several rags-to-riches stories of successful entrepreneurs. Further, the creation of new firms does not automatically benefit investors that hold broad market indices. These indices are composed of existing firms, which sometimes are displaced by new firm entry due to creative destruction. This distinction between new and existing firms introduces a wedge between the growth rate of aggregate dividends and that of dividends per share. To illustrate the contribution of new firm entry to aggregate dividend growth, Figure 1 shows that even though aggregate dividends and aggregate consumption share a common trend, dividends-per-share of the S&P 500 follow a markedly slower growth path. The difference in growth rates between aggregate dividends and dividends per share is approximately 2% per year. This discrepancy can be largely attributed to the dilution effect arising every time new companies enter the index.[1]

In this paper, we exploit fluctuations in the gap between aggregate dividends and dividends per share to identify a 'displacement' shock — that is, a shock that reallocates profits from existing firms to new firms. We show that the identified displacement shock is negatively correlated with returns to the market portfolio and the returns of value strategies. Further, we show that this displacement shock carries a negative risk premium: firms that have higher than average exposures to the shock earn lower than average returns. Last, we show that this displacement shock can account for a substantial fraction — approximately one-third — of the equity risk premium.

We motivate the empirical exercise using a minimal extension to the standard endowment economy model that allows for new firm entry and incomplete markets. The key feature of the model is that the ownership of new firms is randomly allocated to a (small) subset of the population. Importantly, investors cannot sell claims against their future endowment of new firms. As a result, shocks to the relative profitability of new firms lead to the redistribution of wealth from the owners of existing firms to the new entrepreneurs. This wealth redistribution increases the cross-sectional dispersion of consumption growth — most households suffer small losses while a lucky few experience large wealth increases. Since marginal utility is convex, the displacement shock leads to increases in the stochastic factor, or equivalently,

[1] To illustrate this, the line labeled 'Index Dividends + New Cap' in Figure 1 plots the logarithm of dividends-per-share but also adds the cumulative (log) increments in the shares of the index that are due to the introduction of new firms. With this addition, the resulting series co-trends with both consumption and aggregate NIPA dividends.

Figure 1: 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) increments in the shares of the index that are due to the introduction of new firms. Sources: R. Shiller's website, FRED, Personal Dividend Income series, and CRSPSift.

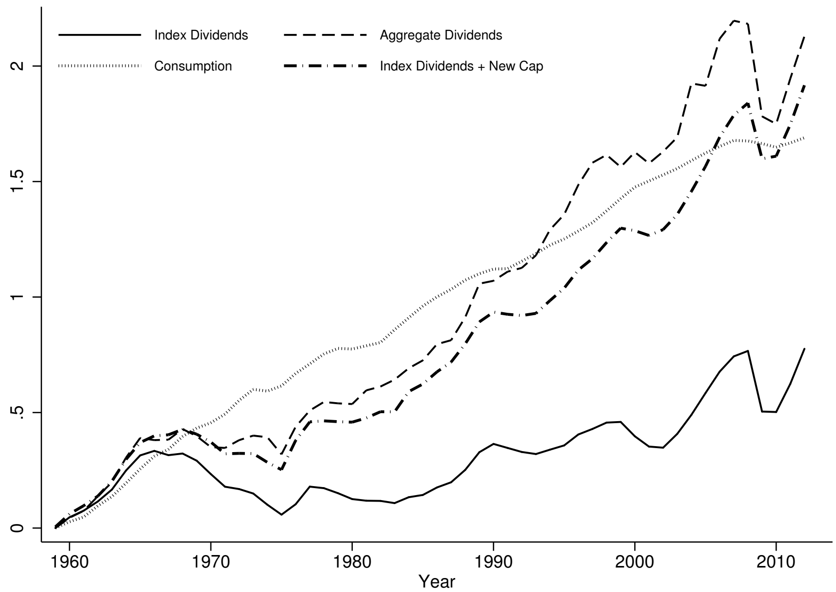

carries a negative risk premium.

Our theoretical model suggests that the displacement shock is closely related to fluctuations in the index divisor — the number of shares of the market portfolio. Using the model as a guide, we estimate a vector error-correction model to decompose the variation in dividends-per-share into 'displacement shocks' — that is, shocks to dividends-per-share resulting from changes in the index — and 'neutral shocks' — that is, shocks that affect dividends-per-share but do not affect the constitution of the index. The advantage of the vector-error correction model is that it can explicitly allow for stochastic delays in the introduction of firms into the market index. Consistent with the theory, the identified displacement shocks have a negative and non-trivial impact on the index dividends-per-share, as well as on the returns of the market index and value strategies. These results hold regardless of whether we focus on the entire CRSP universe or on the S&P 500 index.

We next estimate the risk premium associated with the identified displacement shock. We form portfolios of firms based on estimated firms' betas with the displacement shock. The difference between the average returns to the top and the bottom decile portfolios is approximately -3.5% per year and is statistically significant. Differences in average returns across decile portfolios are fully accounted by differences in risk exposures to the long-short portfolio. Fama-MacBeth analysis provides similar results in that differences in displacement risk are associated with differences in average returns. In sum, the estimated risk premium associated with displacement shocks is negative and substantial: a pure bet on the displacement shock has a Sharpe Ratio that ranges from -0.76 to -1.45 across specifications, though these numbers are not very precisely estimated (t-statistics range from 1.85 to 2.60). This negative risk premium implies that the displacement shock contributes positively to the equity risk premium — since it is associated with market declines.

In the remaining part of the paper, we quantify the contribution of displacement risk to the equity risk premium. We proceed along two fronts. First, we compute the risk premium of a fictitious market index that has no exposure to the displacement shock. We find that this fictitious portfolio carries a risk premium that is approximately two thirds of the equity premium of the market portfolio– implying that roughly one-third of the equity premium in the 1966 to 2012 period is compensation for displacement risk. Second, we present two calibrations of the model that explore the extent to which it can generate a realistic equity premium using conventional parametrization. The first calibration uses the baseline model that features minimal deviations from the endowment economy benchmark. Perhaps surprisingly, this simple model, in which the displacement shock is the only source of aggregate uncertainty, can generate an equity premium of approximately 2% using power utility and a coefficient of relative risk aversion of 10. The second calibration considers recursive preferences, and results in even higher equity premia for even the most conservative assumptions on the impact of displacement risk on dividends-per-share.

In sum, our results demonstrate the importance of recognizing the distinction between aggregate dividends and dividends per share. To the extent to which new firm creation is an important source of profit reallocation, it implies that holders of the market portfolio are exposed to additional risks beyond aggregate economic growth. These additional risks increase the equity premium relative to the benchmark model with a representative agent and and endowment economy. Our paper thus contributes to several strands of the literature.

An extensive literature documents a significant impact of technological progress embodied in new capital vintages on economic growth and fluctuations (see, e.g. Solow, 1960; Greenwood, Hercowitz, and Krusell, 1997; Fisher, 2006). Further, there is significant micro-level evidence documenting vintage effects in the productivity of manufacturing plants. Specifically, Jensen, McGuckin, and Stiroh (2001) find that the 1992 cohort of new plants was 51% more productive than the 1967 cohort in their respective entry years. This difference persists even after controlling for industry-wide factors and input differences. Technological breakthroughs naturally favor new firms at the expense of incumbents, since new entrants have the highest incentives to implement new technologies. Along these lines, Greenwood and Jovanovic (2001) and Hobijn and Jovanovic (2001) provide evidence suggesting that the introduction of IT favored new entrants at the expense of incumbent firms in the early 1970s. We add to this literature by proposing a new measure of the degree of displacement faced by existing firms.

Our theoretical model is closely related to Gârleanu, Kogan, and Panageas (2012) and Kogan, Papanikolaou, and Stoffman (2015), who study the pricing of embodied shocks under incomplete markets. We innovate relative to both papers by proposing a simple model that leads to an attractive empirical measure of displacement affecting existing firms, and show that it carries a significant and negative risk premium. Furthermore, by containing only the minimal deviations from a Lucas-tree economy, our model illustrates analytically how displacement risk enters the stochastic discount factor of financial market participants in a manner similar to Constantinides and Duffie (1996). Especially compared to Gârleanu et al. (2012) our model helps illuminate that displacement risk is priced independent of whether the source of market incompleteness is inter- or intra-generational lack of risk sharing. Accordingly, focusing solely on measures that are meant to capture lack of intergenerational risk sharing (as Gârleanu et al. (2012) do) is likely understanding the impact of displacement.

Our modeling strategy is related to Menzly, Santos, and Veronesi (2004) who model directly firm cashflows as shares of aggregate output; we extend their model to allow for new firm creation and vintage effects. Lastly, our work contributes to the voluminous literature that links economic properties of firms to their risk premia.[2]

An important assumption in our model is that the economic rents that are created by new enterprises are not fully shared with stock market participants. This assumption is in line with Hart and Moore (1994), who show that the inalienability of human capital limits the amount of external finance that can be raised by new ventures. Along these lines, Bolton, Wang, and Yang (2015) characterize a dynamic optimal contract between a risk averse entrepreneur with risky inalienable human capital, and outside investors, and show that the optimal contract leaves the entrepreneur with a significant fraction of the upside gains.

Our model delivers predictions that are consistent with several existing studies. Brav, Constantinides, and Geczy (2002) document that the equity premium and the premium of

[2]An incomplete list includes Jermann (1998); Gomes, Kogan, and Zhang (2003); Kogan (2004); Carlson, Fisher, and Giammarino (2004); Zhang (2005); Lettau and Wachter (2007); Kaltenbrunner and Lochstoer (2010); Santos and Veronesi (2010); Gourio (2011, 2012); Ai, Croce, and Li (2013); Belo, Lin, and Bazdresch (2014); Kogan and Papanikolaou (2014); Kung and Schmid (2015); Ai and Kiku (2013); Favilukis and Lin (2013); Croce (2014); Gârleanu, Panageas, and Yu (2012).

value stocks over growth stocks is consistent with a stochastic discount factor calculated as the weighted average of individual households marginal rate of substitution. Johnson (2012) finds that financial assets, such as growth stocks, that hedge against increases in inequality earn lower risk premia. Loualiche (2013) shows that firms in industries more vulnerable to new entry earn higher risk premia.

Last,we are not the first to point out the difference between the dividend growth of an index and the growth in dividends per share. For instance, Pastor and Veronesi (2006) recognize this distinction when reconciling the higher than average future profitability of firms listed in the Nasdaq index at a given point in time with the low (future) profitability of the index itself. We use the distinction between aggregate dividends and dividends per share to create a new index of displacement affecting existing firms and to study its asset pricing implications.

2 The Model

To expedite the presentation of the main results, we start with an intentionally stylized, minimal extension of a standard endowment economy. Section 2.1 introduces the basic model, where we deviate from the endowment economy benchmark by allowing for the creation of new firms. In Section 2.2 we introduce the notion of displacement risk, and show that this risk is priced if the market for ownership claims on these new firms is incomplete. We conclude this section with a discussion of several extensions to the basic model, which are considered in the appendix.

2.1 Firms

Time is discrete and indexed by t. There is an expanding measure of firms, indexed by (i, s) where s denotes the date at which the firm is created and $i \in [0, 1]$ denotes the index of the firm within its cohort. Firms employ labor using a decreasing-returns-to-scale technology. A firm (i, s) produces output at time t according to

$$y_{t,s}^{(i)} = a_{t,s}^{(i)} Y_t, \tag{1}$$

where $a_{t,s}^{(i)} > 0$ captures the fraction of aggregate output accruing to firm i (i.e., $\sum_{s\leq t}\int_{i\in[0,1]} a_{t,s}^{(i)} = 1$). The aggregate output process evolves according to

$$\Delta \log Y_{t+1} = \mu + \sigma\, \varepsilon_{t+1}, \tag{2}$$

where ε_t are i.i.d. according to some known distribution.

We depart from the standard setting by introducing a distinction between new and old firms. Each period a new set of firms arrive exogenously. These new firms, indexed by $i \in [0,1]$, are heterogeneous in their productivity. Specifically, the productivity of a newly arriving firm i satisfies

$$a_{t,t}^{(i)} = \left(1 - e^{-u_t}\right) dL_t^i, \tag{3}$$

where u_t is a random, non-negative, cohort-specific component, affecting all firms born at the same time. The component L_t^i denotes a cross-sectional measure and its increment dL_t^i is a random, non-negative, idiosyncratic productivity component, which is drawn at the time of the firm's birth (and remains unchanged thereafter) and satisfies $\int_{i\in[0,1]} dL_t^i = 1$.

The productivity of firms created at earlier times $s < t$ is given by

$$a_{t,s}^{(i)} = a_{s,s}^{(i)} e^{-\sum_{n=s+1}^{t} u_n}. \tag{4}$$

Combining equations (A.11) and (3), we see that the total fraction of output produced by the cohort of firms born at time t is equal to

$$\frac{y_{t,t}}{Y_t} = 1 - e^{-u_t}. \tag{5}$$

Conversely, the fraction of time-t output due to older firms is e^{-u_t}. Thus, the u_t shock leads to displacement of older firms.

An important feature of our model is the distinction between aggregate dividend growth Y_t, and the growth of the dividends accruing to a unit dollar investment in the market index — that is, the growth in the 'dividends-per-share' of the market portfolio. Denoting by $Y_{t,s}^e$ the dividend at time t of a portfolio consisting of all firms that exist at time s, we have

$$\log Y_{t+1,t}^e - \log Y_{t,t}^e = \log Y_{t+1} - \log Y_t - u_{t+1}. \tag{6}$$

Examining equation (6), we see that the displacement shock u_t introduces a discrepancy between the dividend growth of existing firms (the left-hand side) and the aggregate dividend growth, which is given by $\log Y_{t+1} - \log Y_t$.

A straightforward way to express this discrepancy is through the notion of an index divisor. A divisor S_t captures the number of fictitious "shares" of an index, and its growth rate (S_{t+1}/S_t) is determined so that

$$R_{t+1}^{\text{ex}} = \frac{\frac{P_{t+1}}{S_{t+1}}}{\frac{P_t}{S_t}} = \left(\frac{S_t}{S_{t+1}}\right)\left(\frac{P_{t+1}}{P_t}\right), \tag{7}$$

where R_{t+1}^{ex} is the ex-dividend gross return on the market-weighted portfolio of all firms in existence at time t and P_t represents the aggregate value of the market at time t. By construction, the divisor captures the discrepancy between the (ex-dividend) returns of the market portfolio and the proportional increase in total market capitalization.

Equations (6) and (7) together imply that changes in the divisor are closely related to the displacement shock:[3]

$$\log S_{t+1} - \log S_t = u_{t+1}. \tag{8}$$

The reason for the difference between aggregate dividends and dividends-per-share is that aggregate dividends *do not* constitute the gains of a self-financing strategy. An investor holding the market portfolio needs to pay to acquire the new firms that enter the index. To restore the self-financing nature of the strategy, the investor needs to constantly liquidate some of the shares she holds to purchase shares of the new firms. Since the fraction of the index that this investor holds diminishes over time, the dividend growth of such a self-financing strategy falls short of the growth in aggregate dividends.

This distinction between dividends-per-share and aggregate dividends has two important implications for asset prices. First, the dividends-per-share of an index behave differently from aggregate dividends. Indeed, the two variables are not even co-integrated, since $\log(Y_t/Y_t^e)$

[3]To see this note that

$$\begin{aligned}\log R_{t+1}^{\text{ex}} &= \log Y_{t+1}^e - \log Y_t^e + \log\left(\frac{P_{t+1}^e}{Y_{t+1}^e}\right) - \log\left(\frac{P_t^e}{Y_t^e}\right) \\ &= \log\left(\frac{S_t}{S_{t+1}}\right) + \log\left(\frac{Y_{t+1}}{Y_t}\right) + \log\left(\frac{P_{t+1}}{Y_{t+1}}\right) - \log\left(\frac{P_t}{Y_t}\right),\end{aligned}$$

where the first equation follows from the definition of a gross return and the second equation follows from (7). Using the fact that $\frac{P_t}{Y_t} = \frac{P_t^e}{Y_t^e}$ and equation (6) leads to (8).

behaves like a random walk with drift. Second, and more importantly, the value of the market portfolio does not equal the present value of future aggregate dividends. Instead, the market portfolio equals the present value of the dividends accruing to the firms *currently in existence*, or equivalently, it equals the present value of the dividends-per-share of the market index.

In sum, our stylized model illustrates how displacement shocks lead to a discrepancy between aggregate dividends and dividends-per-share.

2.2 The pricing of displacement risk

Here, we solve for the equilibrium of the model.

2.2.1. Consumers and markets

The economy is populated by infinitely-lived agents, who maximize their expected utility,

$$U_t = E_t \sum_{s=t}^{\infty} \beta^s \frac{C_s^{1-\gamma}}{1-\gamma}, \tag{9}$$

where $\beta \in (0,1)$ is the subjective discount factor. As in the standard Lucas-tree model, consumers can trade equity claims on existing firms and in a riskless, zero-net-supply bond. Moreover, consumers can trade (zero-net-supply) claims to the realization of the shocks u_{t+1} and the innovation to $\log(Y_{t+1})$.

We deviate from the endowment economy benchmark by assuming a market incompleteness. Specifically, at time zero consumers are equally endowed with all firms in existence at that time. However, from that point onward, consumer i receives firm (i,t) at time t, i.e., a new firm with productivity proportional to $a_{t,t}^{(i)}$, which is random. Importantly, a key market is missing: consumers cannot enter contracts that are contingent on the realized value of their future endowments of new firms.

2.2.2. Equilibrium

We make a simplifying assumption to solve the model in closed form. Specifically, we focus on the limiting case in which firm creation generates extreme inequality, in that only a set of measure zero of firms manage to produce non-zero profits; by contrast, the vast majority of new firms are worthless.

Formally, we assume that, for every t, the distribution of idiosyncratic shocks dL_t^i consists exclusively of (random) point masses. Specifically, we assume that L_t is a discrete measure on $[0,1]$, so that it is an increasing right-continuous, left-limits (RCLL) process that is constant on $[0,1]$ except on a countable set, where it is discontinuous. Both the magnitudes of the jumps in L_t, and the locations of the points of discontinuity are random. In words, only a set of measure zero of consumers obtain the profitable new firms. This assumption ensures that when making consumption and savings decisions, households attach zero probability to the event they receive a profitable firm.

The definition of equilibrium is standard. An equilibrium is a set of price processes and consumption and asset allocations such that a) consumers maximize expected utility over consumption and asset choices subject to a dynamic budget constraint, b) goods market clear, and c) all asset markets clear. The next proposition constructs an equilibrium and describes its properties.

Proposition 1 *Let $F_t(i) : [0,1] \to [0,1]$ denote the distribution of consumption across individuals. Then there exists an equilibrium in which*

$$F_{t+1}(i) = e^{-u_{t+1}} F_t(i) + \left(1 - e^{-u_{t+1}}\right) L_t^i. \tag{10}$$

The dynamics of a stochastic discount factor ξ_t that prices all traded assets are given by

$$\frac{\xi_{t+1}}{\xi_t} = \beta \left(\frac{Y_{t+1}}{Y_t}\right)^{-\gamma} e^{\gamma u_{t+1}}. \tag{11}$$

For a heuristic proof of Proposition 1, consider a simplified model where the value of all new firms is equally and randomly allocated to a measure π of the population. At every point in time, all households can be divided into two groups: households that receive profitable new firms — the newly rich (NR) — and those that do not. Agents have a constant consumption to wealth ratio, hence their consumption process is directly linked to the dividends of the firms they own.[4] Hence, the equilibrium stochastic discount factor can be written as

$$\frac{\xi_{t+1}}{\xi_t} = \beta \left(\frac{Y_{t+1}}{Y_t}\right)^{-\gamma} \left((1-\pi) e^{\gamma u_{t+1}} + \pi \left(\frac{1 - e^{-u_{t+1}}}{\pi}\right)^{-\gamma} \right). \tag{12}$$

[4] We prove this fact in the appendix. This result follows from the facts that (i) households have homothetic preferences, (ii) the endowment of new firms represent a permanent change in wealth, and (iii) households have no labor income. These assumptions imply that all households have the same consumption-to-wealth ratio, which is equal to the dividend-price ratio. Therefore, setting one's consumption equal to the dividends of the firms that one owns is actually optimal.

Our assumption of extreme inequality — that is, that L_t^i is comprised of point masses — implies that as $\pi \to 0$, then (12) is equal to (11).

Examining equation (11), we see that incomplete markets introduce a wedge between our stochastic discount factor and the one arising in a standard, Lucas-tree endowment economy. This additional term, given by $e^{\gamma\, u_{t+1}}$ adjusts for the fact that almost all consumers do not consume the aggregate endowment; a set of measure zero (NR) consumes a non-trivial amount. Households differ in terms of their wealth — and hence their marginal utility — because they cannot share the idiosyncratic risk associated with the future random endowments of firms. Since marginal utility is a convex function of consumption, the average marginal utility is higher than that of an agent who would be consuming aggregate consumption.[5]

2.3 Discussion of the Modeling Assumptions

The model we have presented is intentionally stylized, and only considers an endowment economy. The appendix contains numerous extensions, where we model the production process in greater detail. We summarize these extensions here and provide the key results and intuitions.

Our stylized model abstracts from labor income. Hence, one potential concern is that labor income can serve as a hedge against displacement shocks. However, displacement shocks are also likely to affect labor income. In fact, recent evidence shows that the job market has become increasingly polarized, likely in response to technological advances (Autor, Katz, and Kearney, 2006). To illustrate the interaction between labor income and displacement shocks, we extend the model to allow for labor in the production function and replace the infinitely-lived-agent setting with an overlapping-generations setting. We show that if the human capital of younger workers is better suited to the innovations introduced by new firms, then the stochastic discount factor exhibits displacement risk. Intuitively, the lack of inter-generational risk sharing between cohorts of workers with different human capital and the lack of intra-generational risk sharing for allocations of new firms have similar effects.[6] In fact, if workers born at time s can only provide their labor to firms born at time s — that is,

[5] In this respect, our model is similar to Constantinides and Duffie (1996). However, the two models rely on quite different endowment specifications. Constantinides and Duffie (1996) assume that agents experience permanent shocks to their endowment of consumption, while in our model the increased dispersion in consumption growth derives from new firm creation.

[6] An interesting aspect of allowing for labor income is that in general consumers will not find it optimal to consume their endowments. Asset markets play a non-trivial role in allocating risks. An interesting theoretical implication of this fact is that the stochastic discount factor reflects not only current displacement shocks, but also anticipations of future displacement shocks, similar to models of long-run-risks.

only the younger agents can operate the latest technologies — then the model of the previous section is unaltered.

In our simplified model we have also abstracted away from physical capital. A potential concern is that if households were allowed to invest their savings in the physical capital of existing firms, the displacement risk would be mitigated. However, this is not the case as long as there is some market power or decreasing returns to scale. Allowing for endogenous firm entry leads to similar conclusions, as long as the rents from new firm creation are not perfectly shared across all households.

We conclude this section with two remarks. First, the assumption of extreme inequality is attractive for technical reasons but is not essential for the key intuition. Equation (11) would not hold exactly if an agent's probability of receiving valuable new firms is not zero, but would be a good approximation as long as this probability is small. Kogan et al. (2015) features a similar market incompleteness, but where consumers have a non-zero probability of receiving profitable projects; the results, obtained using numerical solutions, are qualitatively similar.

Second, we did not make any assumption on the correlation between the displacement shock u_t and the shock affecting output ϵ_t. The reason is that equation (11) for the stochastic discount factor holds true irrespective of this correlation. To see this, let $\epsilon_t = \delta u_t + \epsilon_t^{\perp}$, where $\epsilon_t^{\perp}$ orthogonal to u_t. An equivalent way to express (11) is

$$\log \frac{\xi_{t+1}}{\xi_t} = \log \beta - \gamma \epsilon_{t+1}^{\perp} + \gamma(1-\delta) u_{t+1}. \tag{13}$$

Here, the shock $\epsilon_t^{\perp}$ can be interpreted as a 'level' shift in the dividends of all firms (existing and new), which is orthogonal to redistribution in the relative market valuations of the two types of firms.

Equation (11) shows that displacement risk is always priced when used in conjunction with aggregate consumption growth, since it helps 'clean up' the variation in aggregate consumption growth that does not accrue to the marginal agent. Equation (13) shows that displacement risk is a source of risk, even when viewed in isolation from other disturbances, as long as $\delta < 1$. Indeed, $\delta < 1$ is a standard implication of endogenous growth models, tightly linked with the 'creative destruction' that occurs in models of expanding varieties of goods. In these models δ reflects the magnitude of the markup and is usually calibrated to be a number around 0.2.[7] Last, an additional reason why displacement risk is likely to

[7]See, e.g., Gârleanu et al. (2012).

carry a negative risk premium is preferences for relative consumption. Indeed, if agents have preferences of the form $c_t^\eta \left(c_t/C_t\right)^{1-\eta}$, where c_t is an agent's own consumption, C_t is aggregate consumption, and $\eta \in [0,1]$ is a weighting factor, then the coefficient on the displacement shock would equal $1-\delta\eta$.

3 Measuring displacement shocks

We begin our analysis by describing how we can extract a time-series for the displacement shock using movements in the divisor and aggregate dividends. Conceptually, our analysis is based on equation (8), which suggests a straightforward way to measure displacement shocks as changes to the index divisor. In doing so, however, we face three challenges. First, our model implicitly assumes that all firms are traded in the stock market immediately after their introduction; in the data this is not the case. Second, in the data, the divisor is not only affected by new company introductions, but also by corporate payout decisions of existing firms. Last, displacement may take forms that are not captured by new company introductions. For instance, consider a publicly traded firm issuing shares to purchase a non-publicly traded firm. We next describe how we address these issues in detail.

3.1 Choice of the Index

A seemingly important choice in our analysis is the definition of the market portfolio. One possibility is to use the CRSP index as a measure of the market index, and define changes in the divisor as the log-difference of aggregate market capitalization minus the log-difference of the value of the CRSP-index. However, the two significant expansions of the index, in 1962 and 1972, when CRSP starts covering AMEX and NASDAQ stocks, respectively, pose a practical challenge. At these dates the CRSP divisor experiences substantial changes (around 6% and 11%, respectively). These jumps in the divisor in 1962 and 1972 clearly do not represent the creation of new firms, but rather a wave of new listings that reflect past firm entry. Further, these changes in the composition of the sample make it likely that some of the econometric estimates we obtain may not be constant in each sub-sample.

In sum, these jumps introduce outliers, and it is not obvious how to handle them. Another possibility is to consider an alternative definition of the market portfolio that does not suffer from that problem. Specifically, the S&P 500 is an index that is available since 1957 and is professionally maintained to provide 'broad coverage' of the market. Practically, this means that the index has traditionally covered a relatively stable 80% of the market capitalization

Figure 2: Bottom graph: Overall change in the S&P 500 divisor (solid line) and change in the CRSP-divisor (dotted line). Top graph: Change in the S&P 500 divisor (solid line) and change in the CRSP-divisor (dotted line) that is due to new firm entry (first term of equation (14)). To measure new firms in the S&P we use CRSPSift, while to measure new firms in CRSP we use the market valuation of PERMCOs that appear during the respective year.

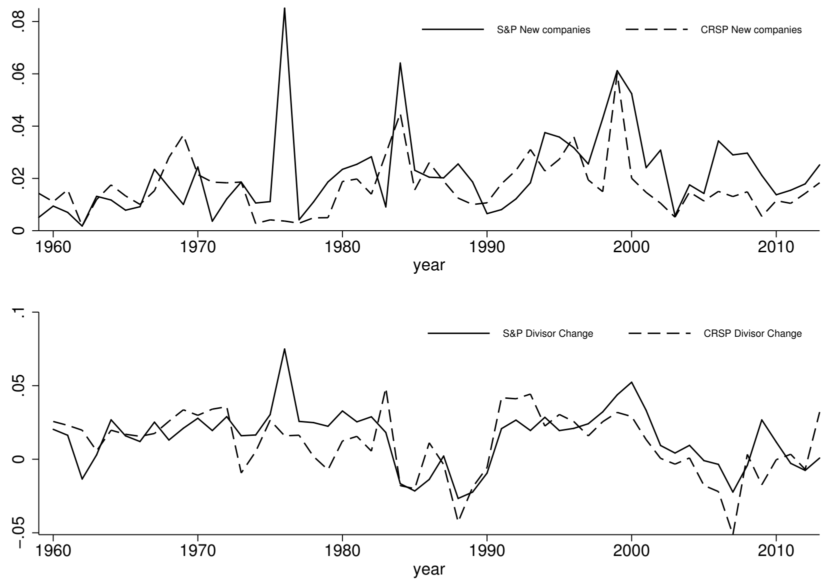

of US markets. In the context of our model, it is straightforward to show that any index that reflects a constant fraction of 'true' market capitalization at any point in time should exhibit the same divisor changes as an ideal index reflecting the entirety of the stock market capitalization. If that fraction is not literally constant, but stationary, then the index should have the same stochastic trend as the ideal index.

The bottom graph of Figure 2 shows divisor changes for the S&P 500 and the CRSP. In these figures we omit the years 1962 and 1972 — otherwise the CRSP outliers in these years would overwhelm the figure. We see that the divisor changes are highly correlated across the two index definitions. This correlation is particularly strong at lower frequencies, and especially in the latter half of the sample.

In our empirical analysis we use both definitions of the market portfolio, and focus on the 1958-2012 period. We report results using both definitions for completeness; however, our results are essentially identical whether we use the S&P 500 throughout or we use the CRSP and simply replace the 1962 and 1972 divisor changes of the CRSP with the respective ones from the S&P 500.

3.2 New firm entry versus issuance by existing firms

In our model, changes in the divisor occur only due to new firm entry. In the data, however, divisor adjustments reflect all non-dividend transfers between existing households and the corporate sector. In particular divisor adjustments also reflect corporate payout decisions by existing firms that take forms other than dividends (repurchases, issuances, etc.).[8]

To gauge the importance of net issuance by existing firms, we decompose changes in the divisor into a part that results from corporate actions of existing firms and a part that is due to the addition of new firms in the index. Specifically, changes to the divisor can be decomposed as

$$\frac{S_{t+1}}{S_t} - 1 = \underbrace{\frac{P_{t+1}^{new}}{P_t^{old}} \frac{1}{R_{t+1}^{ex}}}_{N_t} + \left(\frac{P_{t+1}^{old}}{P_t^{old}} \frac{1}{R_{t+1}^{ex}} - 1 \right), \tag{14}$$

where P_{t+1}^{new} is the market capitalization of firms entering the index in period $t+1$, P_t^{old} is the time-t market capitalization of firms that are in the index at time t and R_{t+1}^{ex} is the gross return on the index excluding dividend payments.

In the data, a substantial portion in the variation of the divisor is due to new firm entry. In particular, the term N_t in (14) captures fluctuations in the divisor due to new firm entry. The term inside brackets captures the change in the divisor due to firms that belong to the index at time t. We plot both series in Figure 3. We see that changes in the divisor due to addition of new firms (the first term in (14)) are very close to the total changes in the divisor in most years — with the exception of the late 1980's and the late 2000's. We emphasize that these conclusions do not depend on the precise definition of index that we use. Indeed, the first term of (14) is strongly correlated — especially at lower frequencies — across the two index definitions that we consider, as the top graph of Figure 2 shows.

In sum, movements in the divisor reflect non-dividend transfers between the existing shareholders and the corporate sector. Our analysis implies that most of the variation in these transfers is due the addition of new firms. However, some changes in the divisor are due to the actions of existing firms, which we wish to isolate. We next propose an econometric methodology to achieve that.

[8]Corporate payout decisions affect the divisor. Changes in denominations — such as stock splits — do not.

Figure 3: Change in the S&P 500 divisor (dotted line) and change in the divisor that is due to new companies (solid line).

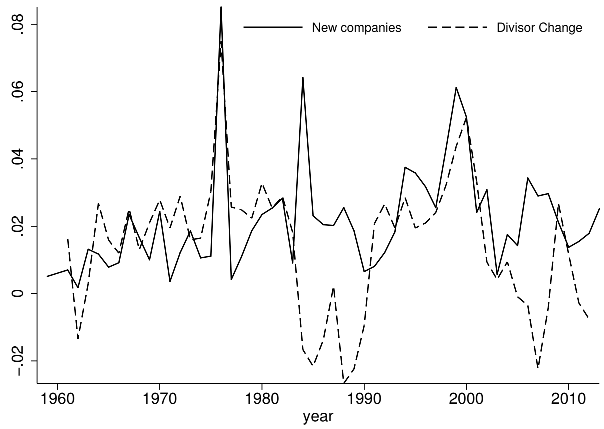

3.3 Identifying the structural shocks

We next consider the identification of the displacement shock from joint variations in the divisor, dividends-per-share, and aggregate dividends. We use two approaches. In our first approach we take the model quite literally. We employ a bivariate Vector Autoregression (VAR) of log dividends-per-share and the log-divisor and infer the 'structural' displacement shock u_t and a neutral dividend shock $\epsilon^{\perp}$, employing a Cholesky decomposition. In simulated data from the model, this approach will correctly identify the displacement shock. In our second approach we recognize explicitly the possibility that variations in the divisor may reflect more than just displacement shocks. We employ a Vector Error Correction Model (VECM) with three variables: log dividends-per-share of the market index, log aggregate NIPA dividends, and log divisor, and use an instrumental variables approach to identify the structural shocks. Throughout, we deflate all dividend series by the Consumer Price Index (CPI).

3.3.1. Identification using a VAR

We start with the VAR analysis. Motivated by the model, we estimate a bivariate VAR of log dividends per share and the log divisor. We extract the (non-structural) residuals of the two series and use a simple Cholesky decomposition to identify the two structural shocks, ordering log dividends per share first. If our model were literally true, then the first shock (the '(divisor-) neutral' shock, labeled 'e' in the figure) would exactly correspond to the shock $\epsilon^{\perp}$ of equation (13) since — by definition — it affects dividends, but not the divisor

Figure 4: Results from a Vector-Autoregression of an identified displacement shock (u) and an identified divisor-neutral shock (e) on dividends-per-share (labeled index dividends) and the divisor of the market index. Dotted lines indicate 95% confidence bands. The left panel presents results using the S&P 500 as the index definition; the right panel presents results using the CRSP value-weighted index. Data period is 1962-2012.

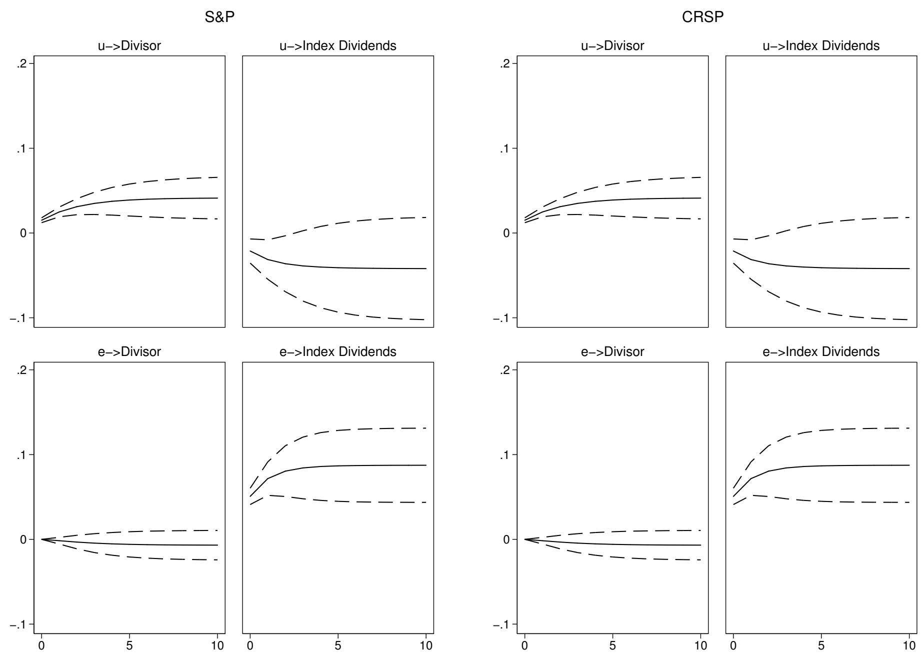

on impact. The second shock (the 'displacement' shock, labeled 'u_t') would then exactly correspond to the shock u_t of (13). We use one lag in the VAR due to the relatively short sample; adding more lags does not significantly impact the results.

We plot the impulse responses to the two structural shocks in Figure 4. We see that both shocks have a significant economic impact on dividends-per-share. Consistent with our model, the neutral shock has a positive impact on the dividends-per-share; by contrast, the identified displacement shock has a negative impact. In terms of magnitudes, a one-standard deviation neutral shock has a long run (permanent) impact of about 8% on dividends, while the divisor shock has a permanent impact of about -4.4% on dividends-per-share. This quantitatively significant *negative* impact on dividends-per-share is consistent with the presence of the displacement shock.

An important advantage of our VAR approach — compared to the seemingly simpler approach of regressing dividends-per-share on the divisor change — is that it allows us to isolate innovations in the two series. This is important because new firms typically enter the index with a lag. For this reason, divisor changes are likely to be a moving average of past displacement shocks. Assuming that this moving-average specification is invertible — a standard assumption in the VAR literature — the VAR analysis identifies correctly the innovations and their impact on the endogenous variables. By contrast, a simple regression would not.

To see why the VAR can still identify displacement shocks even if firms enter the index with a lag, consider the following modification to the baseline model. Firms that are born at time t enter the index at time $t+\tau$, with τ geometrically distributed (with parameter $\hat{\lambda}$) — an assumption that is consistent with the data. In this case, the change in the index divisor S will equal, up to log-linearization,

$$\begin{aligned}\Delta \log S_{t+1} &\simeq \hat{\lambda} \sum_{s=0}^{t+1} \left(1-\hat{\lambda}\right)^{t+1-s} u_s \\ &= \left(1-\hat{\lambda}\right) \Delta \log S_t + \hat{\lambda} u_{t+1}.\end{aligned}$$

Hence, a VAR with one lag should permit correct identification of u_{t+1} despite the lag in introducing past companies into the market.

3.3.2. Identification using a VECM

A limitation of the analysis sofar is that on impact neutral shocks do not affect the divisor. As mentioned, this is a valid assumption if all movements in the divisor are due to displacement shocks. In the data, however, divisor movements may be affected by variations in corporate issuance decisions.

To disentangle the displacement shocks, we exploit two ideas. The first idea is that displacement shocks have –by definition– permanent impact on the divisor. By contrast, if some firms engage in dividend smoothing (i.e., pay the permanent component of their free cash flows as dividends), then at least a fraction of non-dividend corporate payouts are likely to have transitory impact on the divisor. The second idea is to use the information in decomposition (14), which separates divisor movements that are due to new firm entry from divisor movements that are not.

To discuss these issues, we expand our framework of section 2. Specifically, we model aggregate dividends, dividends-per-share, and the divisor as a co-integrated system

$$\Delta y_t = \alpha \, \beta' y_{t-1} + \Gamma_1 \Delta y_{t-1} + v_t, \tag{15}$$

where y_t is a vector containing real-log dividends-per-share, real-log aggregate NIPA dividends, and the log divisor. Here, $\beta' y_{t-1}$ is the co-integrating vector, and v_t denotes residuals with covariance matrix Σ that are uncorrelated across time. (15) is a generalization of the model in section 2.[9]. Our goal is to decompose the residuals v_t of the above system into economically interpretable, "structural" shocks, which satisfy certain properties. In particular, we wish to determine a time series of structural residuals $\eta_t = [u_t, \epsilon^{p\perp}, \epsilon^{n\perp}]$, where we will refer to u_t as a "displacement" shock, $\epsilon^{p\perp}$ as a "permanent neutral" shock and $\epsilon^{n\perp}$ as a transient shock. These shocks are defined by the following properties: the shock $\epsilon^{n\perp}$ is transient, in the sense that its impact on any element of y_t is zero in the long run. The shocks u_t and $\epsilon^{p\perp}$ can have permanent impact on all elements of the vector y_t; the difference between them is that the shock $\epsilon^{p\perp}$ is meant to capture parallel shifts in the (log) dividends of all firms without impacting the relative importance of new vs. old firms, whereas u_t is meant to capture displacement, i.e. shifts in the relative shares of new vs. old wealth.

We define B as a matrix mapping the reduced-form residuals v_t to the structural shocks, $v_t = B\eta_t$. B can be alternatively interpreted as the matrix of impulse-responses to the structural shocks on impact (i.e., at time 0). Our goal is to determine the matrix B, so that we can recover the shocks η_t as $\eta_t = B^{-1} v_t$.

Granger's theorem implies that a decomposition of v_t into two permanent and one transient shock is always possible. Specifically, to identify the transitory disturbances we use Granger's theorem which states that the process y_t can be written as

$$y_t = \Lambda \sum_{s=0}^{t} v_s + y_0^*,$$

where Λ denotes the matrix of cumulative impulse-responses of the (non-structural) residuals v_t. (The matrix Λ is straightforward to compute with any statistical package after estimation of the VECM system). Co-integration implies the following restriction on B

$$\Lambda B = \begin{bmatrix} * & * & 0 \\ * & * & 0 \\ * & * & 0 \end{bmatrix}.$$

In the right-hand side matrix the asterisks denote unrestricted elements. The three zero restrictions on ΛB capture the idea that the temporary shock (which is ordered last without

[9] Specifically, by setting the last column of B to zero, $\alpha\beta = [1 - 11]$ and $B[3,2] = 0$, we recover the model of that section

loss of generality) cannot have any long run impact on the endogenous variables y_t. Granger's theorem implies that L is of reduced rank, so that these three zero restrictions constitute only two linearly independent restrictions on B. The requirement $BB' = \Sigma_v$, where Σ_v is the covariance matrix of the non-structural residuals, adds six more restrictions. Hence, to identify the matrix B, we need one last restriction. Before stating this last restriction, we note that the identification of the transitory shock is not affected by the additional restriction that we need to fully identify B. This last restriction is only useful to decompose the two permanent shocks.

To obtain this last restriction we argue as follows. According to our model, the first component of equation (14), $N_t \equiv \frac{P_{t+1}^{new}}{P_t^{old}} \frac{1}{R_{t+1}^{ex}}$, captures movements in the divisor that are due to the arrival of new firms. To isolate innovations in N_t, we project it on the same right hand side variables as in equation (15) along with its own lag, and refer to the resulting residuals as $u_{1,t}$. In our model $u_{1,t}$ captures the displacement shock in its entirety. However, in the data this is unlikely to be the case for a multitude of reasons: Existing firms may issue equity to purchase non-traded firms. Arguing even more broadly, share issues for executive compensation are also a displacement shock in the spirit of this paper, since they direct company payouts to a small set of the population rather than the average investor. Therefore we cannot use $u_{1,t}$ directly as measure of the displacement shock.

However, we can use it to formulate an orthogonality condition that can help us separate the structural shock $\epsilon_t^{p\perp}$. The reasoning is as follows: By definition, the shock $\epsilon_t^{p\perp}$ is supposed to capture parallel movements in total dividends that do not affect the *relative* valuation of old and new firms at time $t+1$. Inspection of $N_{t+1} = \frac{P_{t+1}^{new}}{P_t^{old}} \frac{1}{R_{t+1}^{ex}} = \frac{P_{t+1}^{new}}{P_{t+1}^{old}}$ shows that N_{t+1} should be unaffected by such shocks. This implies the orthogonality condition $\text{cov}(u_{1,t}, \epsilon_t^p) = 0$ as our last remaining identifying assumption. We wish to emphasize that this orthogonality condition is not due to some "causal" reasoning; it is merely a direct implication of the definitions and properties we wish to assign to the structural shocks and the decomposition we wish to achieve.

We estimate the VECM using Johansen's procedure.[10] We plot the three-by-three impulse response functions in Figure 5 in the appendix. The results of our VECM analysis are similar to those of the VAR model. That is, a one standard deviation in the identified displacement shock is followed by a 3% to 8% drop in dividends per share, depending on the horizon. Hence, displacement shocks have a quantitatively important impact on index dividends. As

[10] As a matter of computational convenience, we note that results are practically identical whether we use our orthogonality restriction or simply perform a Cholesky decomposition of the two permanent shocks, which is readily available option in all statistical packages.

Figure 5: Results of the Vector Error Correction model: Impulse responses of the temporary shock, the permanent "neutral" shock and the displacement shock on real, log-dividends-per-share (labeled index dividends), the log-divisor and log-real, aggregate dividends. Panel A presents results using the S&P 500 as the index definition; Panel B presents results using the CRSP value-weighted index. Data period is 1962-2012.

A. S&P 500

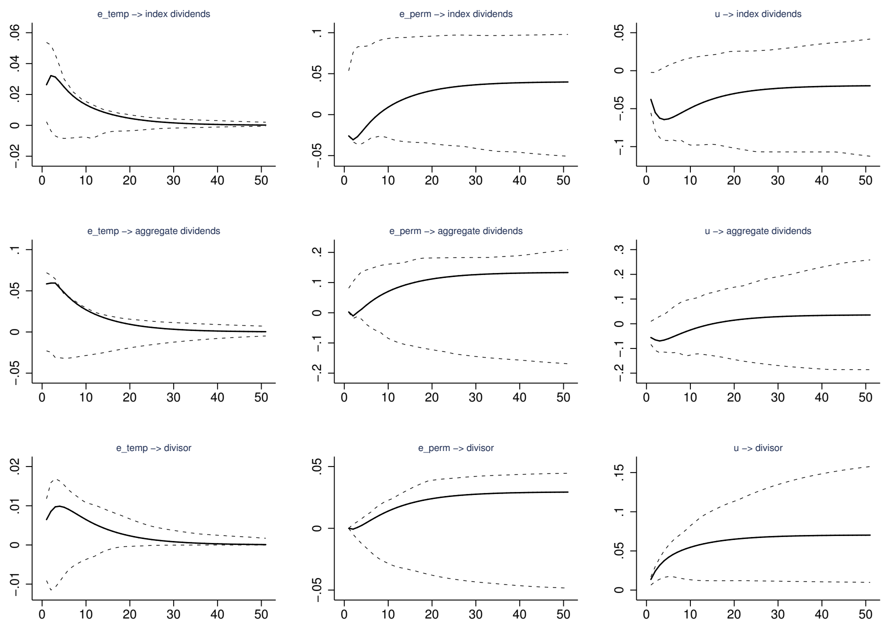

B. CRSP

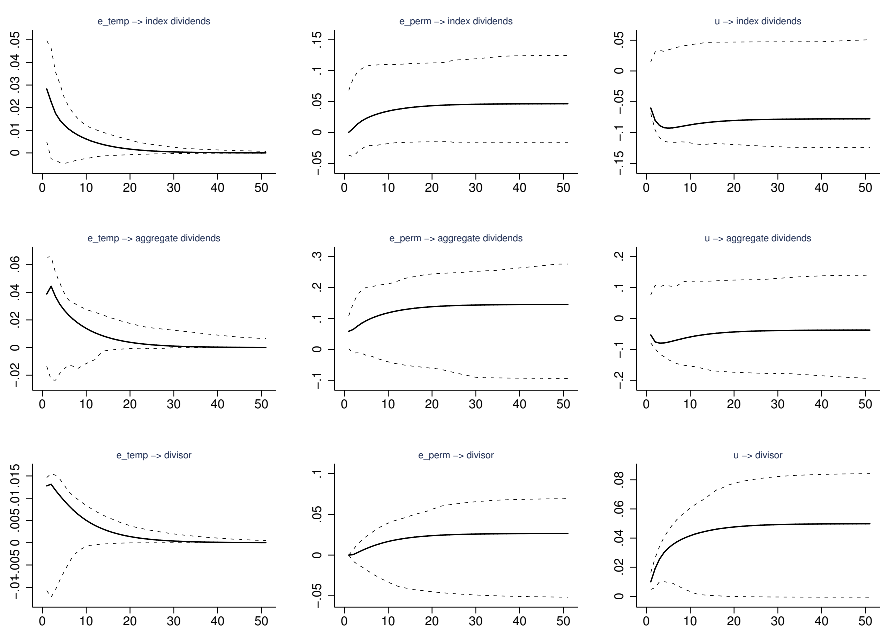

before, the impulse responses are similar no matter how we define the market index (CRSP or S&P). Indeed, when we extract the time series of "displacement" shocks by utilizing either definition of the index, the two series are highly correlated (the first principal component explains 80% of their joint variation). In the next section, we utilize this first principal component, which is likely to contain less measurement error than either series. However, the conclusions are not sensitive to this choice; using either series leads to similar results.

For our purposes, the impulse response functions of either the VAR or the VECM procedure are not the focus of our attention, so we don't discuss them further. Our main goal is to extract the structural shocks, which constitute the input for our asset pricing tests. Before performing these tests, we discuss some possible concerns with our decomposition approach.

3.3.3. Robustness and discussion

We performed several robustness tests to examine the sensitivity of our results. First, we replaced dividends-per-share with earnings-per-share, and aggregate dividends with aggregate earnings in various permutations. Doing so alleviates the concern that firms' payout policy has changed over the last few decades, as firms shifted more to repurchases rather than cash dividends as a form of shareholder payout. Second, rather than imposing that the neutral shock is orthogonal to innovations in the divisor due to new firm entry, we used a simple Cholesky decomposition as in our VECM analysis. The correlations between the identified displacement shocks using any of these procedures and the identified displacement shocks from the procedure we used are fairly high (between 0.73 and 0.99), suggesting that results are robust to these choices.

A typical concern with any analysis that is based on the decomposition of non-structural residuals into structural residuals is whether the interpretation of the structural residuals maps into the correct economic interpretation. One could worry for instance that our instrument is capturing timing shocks, or sentiment shocks that make lots of firms enter the market when discount rates are low (IPO waves). Hence, even if displacement shocks are zero (say u_t is constant), correlated timing choices could be ascribed to variations in displacement. There is a two-fold response to this issue. First, by construction, our identified displacement shocks have permanent negative impact on dividends per-share and the divisor. By contrast variations in timing would only have transitory impact –by definition–, and our procedure removes such transitory components from all series, and in a way that is not dependent on how we disentangle the two permanent shocks. Second, our identified displacement shocks are negatively correlated with excess market returns, which makes it hard to believe that they are sentiment shocks.

The remaining source of concern is that $u_{1,t}$ is positively correlated with the permanent shock $\epsilon_t^{p\perp}$. To start we repeat that such a possibility would go against the definition of the shock $\epsilon_t^{p\perp}$ and the decomposition we wish to achieve. Nonetheless, in appendix D we entertain this possibility, because we wish to investigate the type of errors we would be making in the event of a positive correlation between $\epsilon_t^{p\perp}$ and $u_{1,t}$.

In appendix D we perform the following thought experiment: Instead of using our decomposition approach, we simply postulate one element of the matrix B, so that we can identify B without having to use our orthogonality condition. For a wide range of possible values of the "true" B we adjust the correlation between $u_{1,t}$ and $\epsilon_t^{p\perp}$ so that if we (incorrectly) imposed our orhtogonality condition and performed our shock decomposition, the resulting matrix B would corresponds to the one that we estimate in the data. In other words, we investigate the possibility that our estimated matrix B is the result of a problematic instrument, while the true B is different. For each possible assumption on the true B, we compute the correlations between the displacement shock that would be identified by our procedure and the true u_t (reap. $\epsilon_t^{p\perp}$).

The main conclusion of that exercise is that if our results were due to a problematic instrument, then our inferred displacement shock would exhibit a positive correlation with both the true displacement shock but also with the true permanent neutral shock. This observation is important for our purposes: In the next section we find that our inferred displacement shock has a *negative* price of risk (high u signals "bad times"). Assuming that $\epsilon_t^{p\perp}$ has a positive price of risk (which would hold for any increasing, concave utility function) then a problematic instrument would make us understate the magnitude of displacement risk, since our inferred displacement shock would commingle a (positively priced) neutral shock.

4 Asset pricing implications of displacement shocks

In this section we perform two exercises. First, we investigate whether the displacement shock carries a significant risk premium. Second, we estimate the contribution of the displacement risk to the equity risk premium.

4.1 The pricing of displacement risk

Here, we focus on whether the displacement shock is priced. Specifically, we examine whether stocks with different betas to the displacement shock exhibit different average returns.[11] Specifically, we form 10 portfolios of stocks sorted on estimated betas with respect to the displacement shock at the end of each year.[12] To mitigate the impact of measurement error in betas, we form beta-sorted, equal-weighted portfolios[13] and compute annual returns over the next year.

We report the results in Table 1. Portfolio 1 is the portfolio with the lowest value of pre-ranking beta and portfolio 10 is the one with the highest. Panel A of Table 1 show that average returns are declining as we move from the portfolio with the highest pre-ranking beta to the portfolio with the lowest pre-ranking beta. The difference in average returns between the 10th and the first portfolio is 3.5% per year with a t-statistic of 2.38. Panel B shows the slope coefficient of regressions of the returns of each portfolio on the displacement shock. These post-ranking betas are monotonically declining in absolute value as we move from portfolio 10 to 1. These results are suggestive of a negative risk premium for displacement risk, but not conclusive, since a) post-ranking betas are estimated with measurement error and b) we haven't accounted for the possibility that the results can be explained by other factors. We revisit these issues shortly, when we perform a Fama-Macbeth style test.

Panel C shows that the beta-sorted portfolios have a relatively simple factor structure. To see this, we run a regression of the form

$$R_t^{ei} = \alpha_i + b_i R_t^{em} + s_i SMB_t + h_i HML_t + \gamma_i (p1 - p10) + \varepsilon_t^i, \tag{16}$$

[11]In our baseline model all stocks have the same exposure, but this is only for simplicity. Extending the model to allow different stocks to have different exposures is straightforward.

[12]We first repeat the VECM exercise of the previous section and identify the displacement shock using monthly data on (log, real) dividends-per-share, (log, real) aggregate dividends, and the log divisor utilizing 12 lags. We extract the structural displacement shock from the VECM, labeled u_t. We then estimate betas for all stocks in CRSP with share codes 10 or 11, by regressing the excess returns R_t^i on $u_{t+1\to t+3}$, i.e., on the displacement shock over the next three months (quarter). We adopt this timing convention because dividends are typically paid quarterly and announced in advance of their payment. Excess returns are above the three-month U.S. treasury bill rates as contained in Ken French's data library. Betas are computed in a rolling five-year fashion. To mitigate the effects of illiquidity we require at least fifty non-zero observations. To mitigate the effect of measurement error in the betas, we use Vasicek (1973) to shrink the time-series beta to the cross-sectional mean. Specifically, the Vasicek(1973) procedure computes beta according as a weighted average of the time-series estimate for each stock (β_{TS}) and the cross-sectional mean of betas (β^{XS}): $\beta_i = w_i \beta_i^{TS} + (1 - w_i)\beta^{XS}$, where $w_i = 1 - \sigma_{i,TS}^2 / (\sigma_{i,TS}^2 + \sigma_{XS}^2)$. In our sample the average w_i is about 0.59 and its standard deviation is 0.21. At the end of December of each year, we sort stocks into ten portfolios based on decile breakpoints for β.

[13]As most papers in the literature, we equal-weight the portfolios to mitigate the impact of measurement error in the betas of individual stocks.

Table 1: Decile Portfolios formed on Displacement Beta. Panel A shows annual average returns of ten portfolios formed on beta to the VECM-identified displacement shock. Beta's are calculated using monthly five year rolling windows to the three month forward displacement shock identified in a VECM using monthly data. Portfolios are then sorted on beta at the end of December every calendar year and rebalanced annually. In Panel C, R_t^{em}, S_t and H_t are the Fama-French three factors, and $(p1-10)_t$ is the 1 minus 10 long-short portfolio. Data are annual, 1966 to 2012.

Portfolio	1	2	3	4	5	6	7	8	9	10	10-1
					Panel A: Summary Statistics						
Mean	11.00	11.32	10.61	9.78	10.10	9.73	10.30	8.30	8.83	7.47	-3.53
t-stat	3.20	3.22	3.08	2.91	3.00	2.95	3.09	2.65	2.76	2.35	-2.38
Std Dev	23.60	24.09	23.61	23.06	23.11	22.62	22.89	21.44	21.92	21.80	10.19
					Panel B: $R_t^{ei} = \alpha_i + \gamma_i dshock_t + \varepsilon_t^i$						
u-shock	-0.08	-0.08	-0.07	-0.07	-0.07	-0.07	-0.07	-0.06	-0.07	-0.06	-0.01
t-stat	-3.19	-3.26	-2.65	-2.73	-2.82	-2.83	-2.79	-2.60	-2.99	-2.77	-1.22
R^2	0.19	0.19	0.14	0.14	0.15	0.15	0.15	0.13	0.17	0.15	0.03
					Panel C: $R_t^{ei} = \alpha_i + b_i R_t^{em} + s_i S_t + h_i H_t + \gamma_i D_t + \varepsilon_t^i$						
alpha	-0.01	-0.01	-0.01	-0.02	-0.01	-0.01	0.00	-0.02	-0.01	-0.01	
t-stat	-1.18	-0.72	-1.10	-1.32	-1.19	-0.94	-0.24	-1.26	-0.48	-1.18	
mkt	1.00	0.98	1.04	0.99	1.01	1.05	0.98	0.97	1.00	1.00	
t-stat	17.10	14.64	15.07	15.64	16.58	16.01	15.48	15.14	16.21	17.10	
hml	0.41	0.38	0.47	0.44	0.42	0.53	0.39	0.44	0.36	0.41	
t-stat	5.70	4.54	5.52	5.59	5.55	6.63	5.04	5.63	4.77	5.70	
smb	0.53	0.59	0.48	0.59	0.55	0.47	0.65	0.50	0.53	0.53	
t-stat	7.38	7.14	5.63	7.64	7.35	5.81	8.42	6.42	7.09	7.38	
p1-10	0.79	0.78	0.61	0.43	0.50	0.19	0.24	0.11	0.02	-0.21	
t-stat	8.29	7.10	5.40	4.20	4.99	1.75	2.30	1.06	0.20	-2.22	
R^2	0.93	0.91	0.90	0.91	0.92	0.90	0.91	0.90	0.91	0.92	

where we augment the Fama-French 3 factor model with the long-short portfolio formed by going long the portfolio with the lowest displacement-beta, and short the portfolio with the highest displacement-beta, p1-10.

Examining Panel C of Table 1, we see that the portfolios exhibit relatively similar loadings to the three Fama-French factors, but a clear monotonic loading on the fourth factor. We also performed simple alpha-style regressions. The CAPM alpha of the 1-10 portfolio is 3.5 % with a t-statistic of 1.9, while the 3-Fama French alpha is 3.0 % with a t-statistic of 1.74. The ability of these models to explain some portion of the returns of the 10-1 portfolio is consistent with the model. In our baseline model, for instance, the return of the market portfolio is a linear combination of the displacement shock and the neutral shock. The better performance of the Fama-French model compared to the CAPM is partially due to the fact that the HML factor is negatively correlated with the displacement shock, consistent with (Gârleanu et al., 2012; Kogan and Papanikolaou, 2014; Kogan et al., 2015). Since we wish to test whether the displacement shock is priced in the cross section (and since this shock is not a portfolio return) we next perform a Fama Macbeth-style "horse race" between the displacement shock and the Fama French factors.

Specifically, to confirm that exposures to the displacement shock are priced we start by considering the ten displacement-beta-sorted portfolios as test assets. Motivated by existing work that links displacement shocks to the value premium (Gârleanu et al., 2012; Kogan and Papanikolaou, 2014; Kogan et al., 2015), we also include the 25 Fama-French book-to-market and size portfolios. Here, in addition to the identified displacement shock u_t, we also examine the risk premium associated with the other two structural shocks we recover from the VECM, labeled BC (transitory shock) and PN (permanent, neutral shock). First, we estimate time-series betas by projecting the returns of each asset on the various factors we consider. In a second step, we run a series of cross-sectional regressions of returns on estimated betas and report the associated average coefficients. The standard errors account for measurement error in betas using the Shanken correction.

Table 2 presents the results. We see that higher exposures to the displacement shock are associated with lower average returns. This pattern is robust to controlling for exposures to the other three Fama-French factors. Further, since the displacement shock is normalized to unit standard deviation, the estimated risk premia from the Fama-McBeth regression correspond to the Sharpe Ratio associated with a pure bet on that shock. Depending on the specification, our estimates of the Sharpe Ratio range from -0.76 to -1.45. However, these estimates are fairly imprecise — the t statistics range from 1.93 to 2.60. We conclude that displacement risk is associated with an economically significant risk premium.

Table 2: Fama-Macbeth cross-sectional regressions. Betas are first estimated from the time-series regression $R_t^{ei} = \alpha_i + \beta_i' f_t + \varepsilon_t^i$, f = [mkt smb hml dshock], for each i, using ten portfolios formed on displacement beta and 25 Fama-French portfolios formed on size and book-to-market. Data are annual, 1966-2012. We use raw returns and estimate the risk-free rate in every cross section.

	(1)	(2)	(3)	(4)	(5)	(6)
mkt				0.06		-0.02
				(2.09)		(-.69)
smb					0.03	0.04
					(1.57)	(1.61)
hml			0.06	0.06	0.05	0.05
			(2.61)	(2.58)	(2.46)	(2.44)
BC-shock		-0.19	0.080	-0.070	-0.340	-0.31
		(-.36)	(.17)	(-.17)	(-.86)	(-1.17)
PN-shock		.28	.25	.13	-.13	-.12
		(.39)	(.40)	(.22)	(-.37)	(-.35)
u-shock	-1.39	-1.45	-1.14	-1.06	-.78	-.76
	(-1.93)	(-2.22)	(-2.60)	(-2.10)	(-2.50)	(-2.11)

Figure 6: Displacement Shock vs HML. Figure plots the two-year moving average of identified displacement shock and two-year moving average of $(-1) \times \log(1 + hml)$.

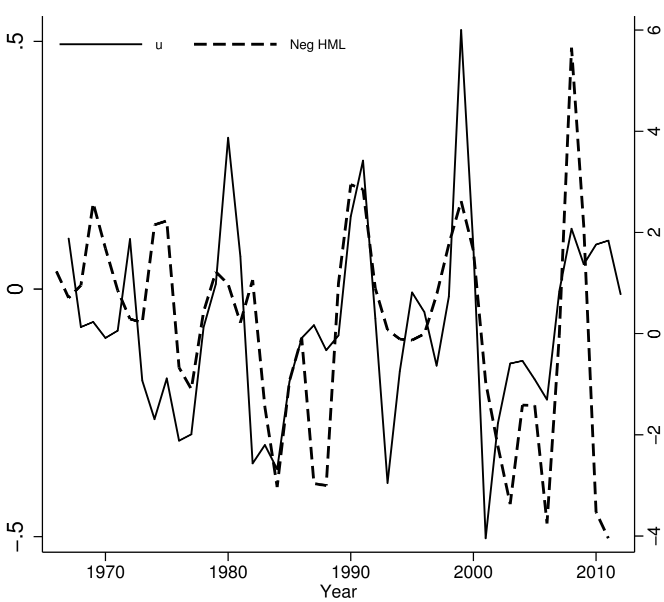

4.2 Displacement risk and the value premium

A natural prediction of models of displacement risk is that growth firms have higher exposure to displacement shocks than value firms (Gârleanu et al., 2012; Kogan and Papanikolaou, 2014; Kogan et al., 2015). The mechanism is that firms in a position to grow ("growth firms") are more likely to benefit from improvements in productivity of new capital vintages than value firms, and thus act as a hedge against displacement shocks.[14] Here, we examine whether this prediction holds using our new measure of displacement risk.

We find that the correlation between our identified displacement shock and the Fama-French HML factor is equal to -20% at the annual level. To adjust for the fact that the portfolios are rebalanced at different dates — Fama-French portfolios are rebalanced at the end of June, while our displacement shock is identified with annual data covering January-December — we also compute correlations at the two-year horizon. In this case, the correlation rises in magnitude to -45%. Figure 6 plots the two-year moving averages of the displacement shock against two-year moving averages of $\log(1 + hml)$. As is evident from the figure, there is a noticeable co-movement between the HML factor and the (absolute value of) the displacement shock.

4.3 Displacement risk and the equity premium

Since the market portfolio is a claim on the *existing* set of firms, displacement risk also affects the equity premium. Here, we quantify the contribution of this risk to the equity premium using a hedging exercise. The goal is to compute the excess return of a portfolio that resembles the market portfolio, except it is "hedged" for displacement risk.

To formalize this notion we start by performing the following regression:

$$R_t^M - (1 + r_t^f) = \alpha + \beta u_t + \eta_t, \tag{17}$$

where $R_t^M - (1 + r_t^f)$ is the excess return on the market portfolio, u_t is the identified displacement shock and η_t are the residuals of the regression. By construction, these residuals correspond to the part of the excess market return that is orthogonal to the displacement factor u_t. We will refer to these residuals as the "hedged payoff".

If u_t was a traded factor (say the payoff of a fictitious security that pays off u_t) then $-\alpha$ would correspond to the average excess return associated with a portfolio that pays off the hedged payoff η_t. Since u_t is not the excess return of a traded factor, we cannot assign such an interpretation to α.

[14] In the context of our model, a new firm (i, s) can be interpreted as a new project, which may be introduced by an existing firm or a new firm. Under such a modified model, growth firms are those with a higher chance of acquiring new projects.

However, we can project η_t on a linear combination of returns from traded zero-investment portfolios $R_t^{e,i}$ that are meant to replicate η_t

$$\eta_t = \alpha_e + \sum_i \beta_i R_t^{e,i}. \tag{18}$$

The coefficients β_i can be interpreted as the weights of the replicating portfolio, while $-\alpha_e$ can be interpreted as the average excess return of the replicating portfolio.[15] For the set of excess returns that are meant to replicate the hedged payoff we choose various combinations of zero-cost portfolios, such as the three Fama French factors, the p1-10 portfolio and the factor-mimicking portfolio implied by the second-stage Fama MacBeth regression (the time-series of the slope coefficients on the displacement shock beta[16]).

Parenthetically, we note that the above approach to "pricing" the hedged payoff is mathematically equivalent to using the ex-post efficient mean-variance portfolio –formed by the test assets– as a proxy for the stochastic discount factor and then imputing the expected excess return of the hedged portfolio as the covariance between η_t and the ex-post mean-variance efficient portfolio.

We report the results of the hedging exercise in the left part of Table 3. Specifically, we report the results from projecting η_t on various combinations of the five test assets. We see that the R^2 from this projection is fairly high, which implies that the fraction of η_t that is not captured by the test assets is quite low. The row labeled "Avg Ret" reports the average return of this market portfolio that is hedged for displacement risk. The difference in the average overall market return and the hedged market return is given in the row "Diff".

We see that the estimated average return of the hedged portfolio is about 4%. Our estimates thus imply that approximately one-third of the overall equity premium is due to the identified displacement shock.

Examining the portfolio weights of this displacement-neutral market portfolio reveals that, consistent with our discussion above, this portfolio overweighs growth firms relative to value firms. Specifically, all of the portfolios place a negative weight weight on the HML factor and a weight less than one on the market (0.89-0.91). Further, when we include excess returns such as the 1-10 portfolio and the factor-mimicking portfolio implied by the Fama-Macbeth procedure, there is an additional tilt towards these portfolios in an effort to

[15] This follows since the mean of η_t is zero by construction, and hence the constant α_e is equal to $-\sum_i \beta_i \bar{R}_t^{e,i}$, where $\bar{R}_t^{e,i}$ is the average excess return of each test asset i.

[16] Note that the time-series of the slope coefficients on the displacement shock beta is a linear combination of the returns of the test assets included in the Fama MacBeth procedure – hence corresponds to the return of a zero cost portfolio, since all test assets in the Fama Macbeth procedure are excess returns.

Table 3: Equity Premium and Displacement Risk. The first panel ($R^m - \beta u$) reports regressions of the residual payoff of the market portfolio after regressing the identified displacement shock on the excess returns of various portfolios. The second panel ([01]) regresses the zero vector on the same set of excess returns, imposing the restrictions that the beta of the resulting portfolio is zero with respect to the identified displacement shock (and unity with respect to the part of the market return that is orthogonal to the identified displacement shock).

	$R^m - \beta u$			[0 1]		
	(1)	(2)	(3)	(1)	(2)	(3)
mkt	0.90	0.93	0.90	0.68	0.95	0.73
	(21.29)	(27.88)	(22.70)	(9.74)	(16.09)	(21.41)
hml	-0.11	-0.07	-0.11	-1.17	-0.18	-0.88
	(-2.19)	(-1.76)	(-2.14)	(-24.27)	(-.91)	(-6.74)
smb	0.00			-0.12		
	(-.05)			(-.46)		
fmimick		0.01			0.03	
		(4.65)			(4.63)	
p1-10			-0.07			-0.66
			(-1.08)			(-2.35)
R^2	0.93	0.96	0.93			
Avg Ret	4.35	3.94	4.13	-1.62	2.32	-2.04
	(5.46)	(6.06)	(5.10)	(-.69)	(.80)	(-.80)

hedge even more of the variation in the identified displacement shock. The tilt on the market to a value less than one is also not surprising. A regression of the market excess return on the identified displacement shock gives an estimate of $-.05$, implying that a positive one-standard deviation shock to the displacement shock moves the stock market down by 5%. In short, the displacement-hedged portfolio tilts away from the market portfolio and towards stocks that have lower exposure to displacement (such as growth stocks) that have lower average returns than the market.

One concern with the method we utilized above is that even though the payoff η_t is orthogonal to u_t, the return of the replicating portfolio obtained through the regression (18) may not be. To account for that, we also utilize an alternative approach to “pricing” the payoff η_t. Specifically, we construct a portfolio that has maximal correlation with η_t, is normalized to have a beta of one to η_t and importantly has a beta of zero to the identified displacement shock. Such a portfolio can be determined in a constrained regression framework, by regressing the zero vector on the excess returns of the test assets, while constraining the coefficients of the regression so that the resulting portfolio has a beta of zero to the displacement shock and a beta of one to η_t.[17] The results are reported on the right-hand panel of table 3. Compared to the left panel, this alternative approach yields even lower values for the average return of the hedged portfolio. The intuition for this fact has to do with how the two approaches treat the part of the displacement shock that is not spanned by traded assets. The approach on the left panel essentially ignores it, i.e., assigns a zero price of risk to the component of displacement risk that is not spanned by traded assets. The approach on the right hand side of the panel assigns the same price of risk as for the part that is spanned. As a result, the approach on the left panel is more conservative, and therefore forms our base-case.

In sum, our results show that approximately one-third of the equity risk premium — approximately 2% — can be attributed to displacement risk. Next, we explore the extent to which these calculations are consistent with a reasonable parametrization of our model.

5 Calibration

We next perform a calibration exercise to investigate the plausibility of the numbers we obtained in the previous section — namely that displacement risk can account for approximately one-third of the equity premium. To do so, we consider two versions of the model: a)

[17] To see this, let β_η denote a column vector containing the regression coefficients of each test asset excess return $R^{i,e}$ on the hedged payoff η_t and let β_u denote the respective vector containing the regression coefficients of each $R^{i,e}$ on the identified shock u_t. Take any portfolio w with the property that $w'\beta_\eta = 1$ and $w'\beta_u = 0$. Then the correlation coefficient between the portfolio excess return $w'R^{i,e}$ and η_t is given by

$$\rho_{w'R^{i,e},\eta_t} = \frac{cov\left(w'R^{i,e},\eta_t\right)}{\sigma\left(w'R^{i,e}\right)\sigma\left(\eta_t\right)} = \frac{\sigma\left(\eta_t\right)}{\sigma\left(w'R^{i,e}\right)},$$

since

$$cov\left(w'R^{i,e},\eta_t\right) = w'\frac{cov\left(R,\eta_t\right)}{var\left(\eta_t\right)}var\left(\eta_t\right) = \left(w'\beta_\eta\right)var\left(\eta_t\right) = var\left(\eta_t\right).$$

Therefore, maximizing the correlation between $w'R^{i,e}$ and η_t amounts to minimizing the variance of $w'R^{i,e}$ subject to $w'\beta_\eta = 1$ and $w'\beta_u = 0$, which can be formulated as the constrained regression problem

$$\min_{w,\alpha}\left(0-\alpha-w'R^{i,e}\right)^2,$$

subject to $w'\beta_\eta = 1$ and $w'\beta_u = 0$.

the simple model of Section 2.2, yielding Proposition 1, and b) a straightforward extension of the model, whereby persistent components to displacement are priced in the spirit of the long run risks literature. The appendix contains a further calibration of the model allowing for more elaborate dynamics of profits and labor income.

We start with the simple model of Section 2.2. We assume that the displacement shock u_{t+1} is i.i.d. and exponentially distributed with scale parameter θ. (We discuss this assumption shortly). To isolate the role of displacement risk, we assume that aggregate consumption is deterministic, that is, $\sigma = 0$. The equity premium (of unlevered equity) is then constant and equals

$$\frac{E_t R_{t+1}}{R^f} - 1 = \frac{\gamma\theta^2}{(1+\theta)(1-\gamma\theta)}, \tag{19}$$

where R_{t+1} is the gross return on equity, and R^f is the gross return on the risk-free asset. [18] Since returns in the data are levered, we use the Modigliani-Miller formula to convert unleveled to levered returns. Specifically, with the historically observed debt-to-equity ratio the formula suggests that levered equity's expected return is 1.6 times that of un-levered equity.

An attractive feature of (19) is that the equity premium depends only on two parameters, γ and θ. We consider two alternatives for calibrating θ. The first determines θ so that a one-standard deviation shock to u_t results in an increase in the divisor corresponding to the estimated response in the data (see subplot labeled "u-¿Divisor" in figure 4). (Matching the response of dividends-per-share to a u shock would lead to similar values for θ). Second, we also use a more direct, and arguably more conservative, approach by estimating θ from data on the share of new company introductions as a fraction of total market capitalization in annual data.

The first method suggests a value of θ around 0.02 (to match the short run response of the VAR) or 0.04 to match the long-run response. Since in this baseline model all shocks are permanent, arguably the latter number is more relevant, since it captures the permanent impact of u_t on dividends-per-share; as we argue shortly, the long-run response is also more relevant if we consider preferences for late resolution of uncertainty/ The second approach implies a similar value for θ. Estimating θ via maximum likelihood yields a point estimate of 0.023 with a 95% confidence band of $[0.018, 0.03]$. Figure 9 in the appendix– presents

[18] Equation expresses the equity premium in "ratio form" $\frac{E_t R_{t+1}}{R^f}$ (rather than the more conventional $E_t R_{t+1} - R^f$) to highlight that it depends exclusively on two parameters γ and θ. Clearly the "ratio form" and the conventional form are equal up to a first order approximation.

Table 4: Levered equity premium as a function of risk aversion (γ) and displacement parameter θ. To relate un-levered to levered equity, we use the Modigliani Miller formula which implies that the levered equity premium is 1.6 times the un-levered equity premium, using the historical leverage ratio in aggregate data.

	$\gamma = 6$	$\gamma = 8$	$\gamma = 10$	$\gamma = 12$
$\theta = 0.02$	0.004	0.006	0.008	0.010
$\theta = 0.025$	0.007	0.010	0.013	0.017
$\theta = 0.03$	0.010	0.015	0.020	0.026
$\theta = 0.035$	0.014	0.021	0.029	0.039
$\theta = 0.04$	0.019	0.029	0.041	0.057

quantile-quantile plots to verify that our assumption of an exponential distribution provides a good fit to the data.[19]

Table 4 reports the levered equity premium implied by our model for various levels of the scale of the displacement shock θ and consumer risk aversion γ. In an economy without any aggregate risk and for values of risk aversion around 10 and a value of θ in the range of 0.025-0.03, the simple version of the model delivers an equity premium of about 1.5-2%. Higher values of θ that would match the volatility of the permanent component of displacement risk (that is, θ in the range $[0.03, 0.04]$) can lead to values of the equity premium in the range of 2% to 5.7% for conventional levels of risk aversion.

The differences between short-run and long-run responses of the divisor to a displacement shock imply higher equity premiums than the ones in table 4. To see this suppose that we model the extent of displacement each year as being autocorrelated. In particular, suppose that u_t is no longer i.i.d., but instead is given by

$$u_{t+1} = (1 - \phi)u_t + \phi z_{t+1}, \tag{20}$$

where z_{t+1} is i.i.d. and exponentially distributed with scale parameter θ. With specification 20, the magnitude of the short run impact of the displacement shock on the divisor is equal to $\phi\theta$, while the long-run response is $\frac{\phi\theta}{1-\phi}$.

[19]Figure 9 shows q-q plots for the empirical quantiles of new company introductions as a fraction of total market capitalization against the respective quantiles of an exponential distribution. The exponential distribution seems to provide a good fit to the data, since the empirical quantiles align well on the theoretical line. To ensure that the results are not driven by unreasonable assumptions on extreme values, we also perform a simple Monte-Carlo exercise: We compare the Monte-Carlo distribution of the maximum value drawn from the estimated exponential distribution against the respective value in the data. We show that the maximum value in the data is well within the 95% range of values suggested by the Monte Carlo simulations.

	$\gamma = 6$	$\gamma = 8$	$\gamma = 10$	$\gamma = 12$
		$\phi = 0.4$		
$\phi\theta = 0.02$	0.011	0.017	0.025	0.036
$\phi\theta = 0.025$	0.018	0.030	0.047	0.077
		$\phi = 0.5$		
$\phi\theta = 0.02$	0.008	0.013	0.018	0.025
$\phi\theta = 0.025$	0.014	0.022	0.032	0.047
		$\phi = 0.6$		
$\phi\theta = 0.02$	0.007	0.010	0.014	0.019
$\phi\theta = 0.025$	0.011	0.017	0.025	0.034

Table 5: Levered equity premium as a function of risk aversion (γ) and displacement parameter θ. To relate un-levered to levered equity, we use the Modigliani Miller formula which implies that the levered equity premium is 1.6 times the un-levered equity premium, using the historical leverage ratio in aggregate data.

Under assumption (20), and the additional assumption that the investor has inter-temporal elasticity of substitution equal to one and risk aversion γ, the unlevered equity premium can be computed in closed form

$$\frac{E_t R_{t+1}}{R^f} - 1 = \frac{(\phi\theta)^2 \left[1 - \frac{(1-\gamma)}{1-\beta(1-\phi)}\right]}{(\phi\theta + 1)\left(1 - \phi\theta\left[1 - \frac{(1-\gamma)}{1-\beta(1-\phi)}\right]\right)}. \tag{21}$$

Using a value of $\beta = 0.95$ (the choice of β has little impact on the results), table 5 provides the equity premium resulting from various assumptions on ϕ, while fixing the short-run response $\phi\theta$ to be either 0.02 or 0.025 to be conservative. We consider a baseline version of $\phi = 0.5$ to approximately match the autoregressive coefficient of first differences in the log-divisor series, and study the sensitivity of this choice by considering $\phi = 0.4$ and $\phi = 0.6$.

Table 5 shows that the equity premiums that result from the calibration are non-trivial (between 1.1% and 4.7% for risk aversion values between 6 and 10) even for the very low degrees of autocorrelation that we assume (compared, e.g., to the long run risks literature). We would also like to underscore that this equity premium arises in isolation, i.e. in the absence of aggregate risk.

The calibrations we have performed sofar assume that all income is in the form of profits. Introducing labor income does not affect the calibrations, as long as we assume a) an overlapping generations structure for the arrival of new workers, and b) that the human capital of workers in cohort s is specific to firms that belong to the cohort s. (See appendix A.2 for details). Under these assumptions, the profits of existing firms and the wages of existing workers have the same dynamics, so that the stochastic discount factor, excess returns etc., in the presence of labor income are unchanged.

For the more general case where labor income and profits can follow different dynamics we refer to the calibration in Appendix C. In that calibration we also consider model implications for the risk free rate and its volatility, the level of the price-dividend ratio, predictability of excess returns etc. Since these extensions are tangential to the main point of the paper, we do not include them here.

In sum, the goal of this section was to perform a simple "back of the envelope" calculation. We intentionally used the most parsimonious version of our model, without including the many extensions we consider in the appendix. The intention was to make the calculations transparent and dependent on a minimal amount of parameters. The main conclusion of the section is that a 2% equity premium (without any aggregate uncertainty) is consistent with a calibration of the model using a risk aversion parameter (between 8 and 10), which lies in the range of parameters that are considered plausible.

6 Conclusion

We build a simple model that captures the distinction, largely overlooked in the asset-pricing literature, between the cash-flow properties of the market portfolio and aggregate dividends. Not only does the model accommodate the significantly higher growth in aggregate dividends than in dividends-per-share, but it also allows for the wedge between the two — predominantly due to displacement risk — to be priced, due to market incompleteness. Motivated by the model, we propose a new measure of displacement risk. We show that this measure carries a significant risk premium: firms with high exposures to displacement risk earn higher returns than firms with low exposures. Based on the estimated risk premium we conclude that displacement risk can account for approximately one-third of the equity premium.

References

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.

Ai, H. and D. Kiku (2013). Growth to value: Option exercise and the cross section of equity returns. *Journal of Financial Economics 107*(2), 325 – 349.

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.

Belo, F., X. Lin, and S. Bazdresch (2014). Labor Hiring, Investment, and Stock Return Predictability in the Cross Section. *Journal of Political Economy 122*(1), 129 – 177.

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.

Brav, A., G. M. Constantinides, and C. C. Geczy (2002). Asset pricing with heterogeneous consumers and limited participation: Empirical evidence. *Journal of Political Economy 110*, 793–824.

Carlson, M., A. Fisher, and R. Giammarino (2004, December). Corporate Investment and Asset Price Dynamics: Implications for the Cross-section of Returns. *Journal of Finance 59*(6), 2577–2603.

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

Croce, M. (2014). Long-run productivity risk. a new hope for production-based asset pricing? *Journal of Monetary Economics 66*.

Favilukis, J. and X. Lin (2013). Long run productivity risk and aggregate investment. *Journal of Monetary Economics 60*(6), 737–751.

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

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

Gârleanu, N. and S. Panageas (2015). Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing. *Journal of Political Economy 123*(3), pp. 670–685.

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

Gomes, J., L. Kogan, and L. Zhang (2003, August). Equilibrium Cross Section of Returns. *Journal of Political Economy 111*(4), 693–732.

Gourio, F. (2011, March). Putty-clay technology and stock market volatility. *Journal of Monetary Economics 58*(2), 117–131.

Gourio, F. (2012, October). Disaster Risk and Business Cycles. *American Economic Review 102*(6), 2734–66.

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

Greenwood, J. and B. Jovanovic (2001). The it revolution and the stock market. *American Economic Review 89*(2), 116–122.

Haltiwanger, J. (2012, April). *Job Creation and Firm Dynamics in the United States*, pp. 17–38. University of Chicago Press.

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.

Hobijn, B. and B. Jovanovic (2001). The information-technology revolution and the stock market: Evidence. *American Economic Review 91*(5), 1203–1220.

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, April). Asset pricing in production economies. *Journal of Monetary Economics 41*(2), 257–275.

Johnson, T. C. (2012). Inequality risk premia. *Journal of Monetary Economics 59*(6), 565–580.

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

Kogan, L. (2004). Asset prices and real investment. *Journal of Financial Economics 73*(3), 411–431.

Kogan, L. and D. Papanikolaou (2014). Growth opportunities, technology shocks, and asset prices. *The Journal of Finance 69*(2), 675–718.

Kogan, L., D. Papanikolaou, and N. Stoffman (2015). Winners and losers: Creative destruction and the stock market. Working Paper 18671, National Bureau of Economic Research.

Kung, H. and L. Schmid (2015). Innovation, growth, and asset prices. *The Journal of Finance 70*(3), 1001–1037.

Lettau, M. and J. A. Wachter (2007). Why is long-horizon equity less risky? a duration-based explanation of the value premium. *The Journal of Finance 62*(1), 55–92.

Loualiche, E. (2013). Asset pricing with entry and imperfect competition. Technical report, Working paper.

Menzly, L., T. Santos, and P. Veronesi (2004, February). Understanding Predictability. *Journal of Political Economy 112*(1), 1–47.

Pastor, L. and P. Veronesi (2006). Was there a nasdaq bubble in the late 1990s? *Journal of Financial Economics 81*(1), 61 – 100.

Santos, T. and P. Veronesi (2010). Habit formation, the cross section of stock returns and the cash flow risk puzzle. *Journal of Financial Economics 98*(2), 385 – 413.

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.

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

A Appendix

A.1 Extensions

Here, we extend the baseline model to address several issues: a) the role of labor income, b) the source of rents accruing to the firm owners in the baseline model, c) the introduction of capital, and d) endogenous entry into the innovation sector. Our goal is to show that all of these realistic extensions do not impact the key qualitative feature of our model that displacement risk is priced.

A.2 Labor and imperfect intergenerational risk sharing

In the real world, production requires labor. The goal of this section is to show how lack of inter-generational risk sharing across worker cohorts implies an additional reason why displacement risk shows up in the stochastic discount factor. The model is similar to Gârleanu et al. (2012) and the reader is referred to that paper for details. We will show a close connection of the results obtained so far (where idiosyncratic risks are imperfectly shared within a generation) to a model where risk is imperfectly shared across generations.

We wish to highlight two results in this section: First, that under specific assumptions on labor dynamics, we recover exactly the stochastic discount factor (11). Second, more generally, we show a result of theoretical interest: In the presence of labor income, the stochastic discount factor "prices" the persistence of displacement shocks, even though agents have power utilities.

We start by assuming that output is given by

$$y_{t,s}^{(i)} = A_t \left(a_{t,s}^{(i)}\right)^{\eta} \left(l_{t,s}^{(i)}\right)^{1-\eta}, \tag{A.1}$$

where $\eta < 1$and $a_{t,s}^{(i)}$ is as in the baseline version and $l_{t,s}$ are "efficiency units" of labor. Additionally, we drop the assumption that consumers are infinitely lived. Instead, we assume that consumers die with probability λ and a new cohort of consumers of mass λ arrives every period. Consumers arrive in life endowed with efficiency units of labor, but no endowment of firms. We also make the (standard) assumptions in the overlapping-generations-literature that a) there exists a competitive market for annuities, and b) investors maximize their expected life-time utility, but have no bequest motives.

We introduce aging and displacement effects in labor earnings. In contrast to the baseline model, labor is not a homogenous service; instead, the units of labor that enter the production function of firms are measured in terms of a composite service, which is a Cobb-Douglas aggregator of the labor efficiency units provided by workers belonging to different cohorts. Specifically, one unit of labor l_t is given by

$$l_t = \prod_{s=0}^{t} h_{t,s}^{q_{t,s}}, \tag{A.2}$$

where $h_{t,s}$ is the number of hours supplied by workers born at time s, and $q_{t,s}$ is a weighting function satisfying $\sum_{s=0}^{t} q_{t,s} = 1$. Due to the Cobb-Douglas assumption, the wage income of cohort s is given by

$$w_{t,s} = (1-\eta)\, Y_t\, q_{t,s}. \tag{A.3}$$

Thus, we can interpret $q_{t,s}$ as the fraction of labor income going to the worker-cohort s.

We allow for the displacement of human capital similarly to the displacement of firms. Specifically, we assume that the shock u_t also affects the fraction of income accruing to newly arriving, versus existing, workers,

$$\begin{aligned} q_{t,t} &= 1-(1-\delta)\, e^{-\psi u_t}, \\ q_{t,s} &= q_{s,s}\,(1-\delta)^{t-s}\, e^{-\psi\left(\sum_{n=s+1}^{t} u_n\right)}. \end{aligned}$$

Here, the constant ψ captures the exposure of labor to the shock u_t, while δ captures the depreciation of labor income due to aging. The following proposition shows the impact of the displacement shock u_t on the stochastic discount factor.

Proposition 2 *Let*

$$\chi_t \equiv E_t \sum_{s=t}^{\infty} (1-\lambda)^{s-t}\, \frac{\xi_s}{\xi_t} \frac{c_s^{(i)}}{c_t^{(i)}}, \tag{A.4}$$

$$\phi_t^c \equiv E_t \sum_{s=t}^{\infty} \frac{\xi_s}{\xi_t} e^{-\sum_{v=t+1}^{s} u_v} \left(\frac{Y_s}{Y_t}\right) \tag{A.5}$$

$$\phi_t^l \equiv E_t \sum_{s=t}^{\infty} \frac{\xi_s}{\xi_t} (1-\delta)^{s-t} e^{-\psi \sum_{v=t+1}^{s} u_v} \left(\frac{Y_s}{Y_t}\right). \tag{A.6}$$

Then,

$$\left(\frac{\xi_{t+1}}{\xi_t}\right) = \beta \left(\frac{Y_{t+1}}{Y_t} \frac{1}{1-\lambda}\right)^{-\gamma} \left[\eta \frac{\phi_{t+1}^c}{\chi_{t+1}} e^{-u_{t+1}} + (1-\eta)\, \frac{\phi_{t+1}^l}{\chi_{t+1}}\, (1-\delta)\, e^{-\psi u_{t+1}}\right]^{-\gamma}. \tag{A.7}$$

Comparing (A.7) with (11) shows that the lack of intra-generational risk sharing (inability of existing agents to share the endowment risk associated with new firms) and the lack of inter-generational risk sharing (inability to trade with the newly arriving cohort of workers before their birth) have similar effects on the stochastic discount factor.

A novel implication of Proposition 2 is that shocks to the *distribution of future* random variables are priced. Thus, not only does the SDF ξ_t depend on the redistribution of u_t in addition to aggregate consumption Y_t, it also depends on the distribution of future shocks u_{t+s}.

As a concrete illustration of how future shocks are priced, suppose that u_{t+1} is drawn from one of two distributions, F_0 and F_1, according to a Markov regime-switching process $s_t \in \{0,1\}$ with given transition matrix. The Markov property of s_t implies that the valuation ratios χ_t, ϕ_t^c, and ϕ_t^l are exclusively functions of s_t. Equation (A.7) implies that, as long as the ratios ϕ^c/χ and ϕ^l/χ are not constant,

$$E_t\left(\frac{\xi_{t+1}}{\xi_t}s_{t+1}\right) \neq E_t\left(\frac{\xi_{t+1}}{\xi_t}\right)E_t\left(s_{t+1}\right) = \frac{\Pr\left(s_{t+1}=1|s_t\right)}{1+r_t^f}, \tag{A.8}$$

where r_t^f is the one-period risk-free rate. In words, innovations to s_{t+1} command a risk-premium.

The pricing of variables other than one-step-ahead consumption growth might appear puzzling in a model with expected utility. This effect arises because dividends and labor income have (potentially) different exposures to the displacement shock. The growth rate in the consumption of the marginal agents reflects the proportion of wealth owned next period by current agents who survive until then but do not become new rich — in equation (A.7), the term in square brackets. This proportion is a weighted average of the displacement of dividend, respectively labor income, and the relative weight is state dependent as long as these two claims depend differently on the displacement shock. By contrast, in the special case in which labor income and dividends are identical in terms of displacement risk ($\delta = 0$ and $\psi = 1$), the last term in (A.7) collapses to $e^{\gamma u_{t+1}}$, as in the baseline model.

Remark 1 *We note that there is another special case in which the stochastic discount factor* (A.7) *becomes proportional to* (11)*. This is when instead of assuming* (A.2)*, we assume instead that* $l_{t,s} = h_{t,s}$*, i.e., workers of vintage* s *can only work in firms of vintage* s*.*

A.3 Capital and monopolistic rents

In this subsection we investigate the version of the model where labor rather than capital is used in production. For some of our results, we assume non-zero economic profits. We start by showing how either imperfect competition or decreasing returns to scale will produce such an outcome. We then present the full model with capital.

A.3.1. Monopolistic rents

First, we illustrate how monopolistic competition can lead to rents accruing to the firm owners in a manner formally equivalent to assuming decreasing returns to scale. This section is based on arguments from Romer(1986,1990) and Gârleanu et al. (2012). Hence, it is intentionally brief and the reader is referred to the aforementioned papers for details.

Assume that the production of the final consumption good requires intermediate goods, which are supplied by monopolistic competitors. Specifically, the final good is produced by competitive firms according to the following production technology

$$Y_t = \sum_{i\in I_t} x_{i,t}^{\delta}\omega_{i,t}^{1-\delta}, \tag{A.9}$$

where I_t is the set of all firms in existence at time t, $x_{i,t}$ is the intermediate good produced by firm i, and $\omega_{i,t}$ is a measure capturing the relative importance of each firm in the index.

Firms produce intermediate goods using a constant-returns to scale technology, $x_{i,t} = k_{i,t}$, where $k_{i,t}$ is the capital good used by firm i.

Maximizing the profits of the final-goods-producing firm leads to the familiar demand function for the intermediate good $x_{i,t} = \left(\frac{p_{i,t}}{\delta \omega_{i,t}^{1-\delta}}\right)^{\frac{1}{\delta-1}}$, where $p_{i,t}$ is the price of intermediate good i. Using the fact that final good firms make zero profits in equilibrium and that the labor market clears, we obtain that the share of profits accruing to firm i equals

$$\frac{p_{i,t}(x_{i,t})x_{i,t}}{\sum_{i \in I_t} p_{i,t}(x_{i,t})x_{i,t}} = \omega_{i,t}. \tag{A.10}$$

We would have reached exactly the same conclusion, if rather than assuming monopolistic competition we had assumed directly that there are only final-goods-firms, which are competitive and rather than using a constant-returns-to-scale production function, they produce according to

$$y_{t,s}^{(i)} = A_t \left(\omega_{i,t}\right)^{\eta} \left(k_{i,t}^{(i)}\right)^{1-\eta}. \tag{A.11}$$

In sum, this section shows the familiar equivalence of a) assuming that firms produce subject to a decreasing-returns-to-scale technology –as in the baseline model and b) assuming that firms produce according to a constant returns to scale production, but have market power. In either case profits will be non-zero

A.3.2. Capital

We reconsider here the baseline model, where we assume that capital is used in production. Our goal is twofold. First, we show that allowing for capital accumulation leads to qualitatively similar results as the baseline model. Second, we show a close relation between our displacement shock and investment-specific shocks that are commonly considered in the literature.

The production function is given by

$$x_{t,s}^{(i)} = \left(a_{t,s}^{(i)}\right)^{\eta} \left(k_{t,s}^{(i)}\right)^{1-\eta}, \tag{A.12}$$

where $k_{t,s}^{(i)}$ is the amount of capital used by firm i. Assuming that production also requires labor is a straightforward extension.

Capital can be created. Each household can forego one unit of consumption to create one unit of capital. Importantly, capital is specific to each vintage: capital that was created for the cohort of firms born at time s cannot be used for firms in cohort $s \neq s'$. As a result, all capital created at time t is invested in new firms. As before, $a_{t,s}$ evolves according to (3)-(4), that is there is constant arrival of new firms that need new capital to operate and they rent it period by period from existing households.

In this version of the model, the shock u_t can be interpreted as a productivity shock that is specific to capital of vintage t. Since our goal is illustrative and we wish to obtain a simple closed-form solution, we assume that $u_t = \overline{u}$ is constant, that is, the extent of displacement each period is deterministic. Last, purely for convenience and ease of exposition, we assume that investors have logarithmic preferences, $U(c) = \log(c)$.

The following proposition characterizes the steady state of this economy.

Proposition 3 *The consumption growth of all agents (other than the set of measure zero agents who receive new firms) is given by*

$$\frac{c_{t+1}}{c_t} = 1 - \frac{(1-\beta)\,\eta\,(a_{t,t})^{\eta}\,\left(\overline{k}\right)^{1-\eta}\left(1 - \frac{e^{-\eta\overline{u}}}{1+r}\right)^{-1}}{\overline{A}\left(\overline{k}\right)^{1-\eta} - \overline{k}} < 1, \tag{A.13}$$

where $\overline{A} \equiv \sum_{s\leq t} (a_{t,s})^{\eta}$ *and* $\bar{k}$ *is a constant (given in the proof) that is equal to the amount of consumption goods investors forego each period to produce new capital goods.*

Proposition 3 illustrates that the consumption growth of the marginal investor in (A.13) is below the aggregate consumption growth rate. The fact that the production function in (A.12) features decreasing returns to capital, while investment and consumption are perfect substitutes at the margin implies that the benefits of new firm creation accrue (at least partly) to new firm owners rather than all households.

Further, the price of installed capital falls with u_t, illustrating the close link with investment-specific technology shocks. Specifically, the price of installed capital of vintage s at time t ($q_{t,s}$) satisfies the following recursion,

$$q_{t+1,s} = q_{t,s}\, e^{-\eta\overline{u}}. \tag{A.14}$$

That is, the shock $\overline{u}$ leads to the economic depreciation of existing capital, illustrating the close relation between the displacement shock and investment-specific shocks.

A.3.3. Endogenous entry

Next, we allow for endogenous entry of firms. Specifically, we now assume that the creation of a new firm requires human capital. That is, a representative "venture capitalist" hires inventors to produce ideas that will lead to new firms. The output share of new firms created at t is given by

$$a_{t,t} = l_t\left(1 - e^{-u_t}\right) \tag{A.15}$$

where l_t is the amount of inventors hired by the venture capitalist. Since (A.15) is linear in l, the venture capitalist makes zero profits. Importantly, innovators are assigned to one project, and only a set of measure zero of these projects are economically viable, as in the baseline model. However, the ones that prove to be viable result in the same (large) profits as in the text.

A key assumption is that once the idea for a new firm proves to be viable, the innovator who came up with the idea needs to be given a fraction ν of the market value of the associated firm, so that he has incentives to develop the project to completion. The innovator can appropriate a fraction ν either because he is essential to the success of the project, or because she can steal the idea and start her own firm. Here, we should emphasize that the term innovator includes not just the entrepreneur who had the idea for the new firm, but also partners in the VC firm who had the talent to discover the entrepreneur as well as early stage employees. In sum, the fraction ν captures the share of the project value that does not accrue to outside investors that buy shares in an IPO.

The following proposition characterizes the stochastic discount factor

Proposition 4 *The stochastic discount factor is given by*

$$\left(\frac{\xi_{t+1}}{\xi_t}\right) = \beta \left(\frac{Y_{t+1}}{Y_t}\right)^{-\gamma} \times \left\{1 - \frac{\nu}{\chi} M_{t+1,t+1}\right\}^{-\gamma}, \tag{A.16}$$

where $M_{t,t} = \phi\, a_{t,t}$ is the aggregate market value of projects created at time t, ϕ is a constant, and χ is the consumption-to-wealth ratio, which is constant.

Proposition 4 shows that even if there is endogenous entry, the stochastic discount factor has the same form as in the baseline model. Intuitively, as long as a fraction of innovation is inalianable from the innovators, and the proceeds from innovation are not equally shared in the population, then innovation will make some parts of the population disproportionately rich and the key insights of the paper continue to hold.

B Proofs

Proof of Proposition 1. Let $\mathcal{N}_{t+1}$ denote the set of all indices of agents (measure 1) who receive a worthless firm. Then we have that

$$\frac{\xi_{t+1}}{\xi_t} = \beta E \left(\frac{c_{t+1}^{(i)}}{c_t^{(i)}}\right)^{-\gamma} = \beta^\tau \left(\frac{\int_{i \in \mathcal{N}_{t+1}} dC_{t+1}^{(i)}}{\int_{i \in \mathcal{N}_{t+1}} dC_t^{(i)}}\right)^{-\gamma} \quad \forall \tau \geq 0, \tag{B.1}$$

where the first equation follows from the consumer's Euler equation and the second equation follows from the probability of receiving a valuable firm being zero. Hence a consumer's anticipated consumption growth coincides with the consumption growth of the cohort that does not receive a valuable firm at $t+1$. Market clearing implies that

$$C_{t+1} = \int_{i \in \mathcal{N}_{t+1}} dC_{t+1}^{(i)} + \int_{i \notin \mathcal{N}_{t+1}} dC_{t+1}^{(i)}. \tag{B.2}$$

We next observe that since the product of the two indicater functions we have that

$$\int_{i \in \mathcal{N}_{t+1}} dC_t^{(i)} = \int_{i \in [0,1]} 1_{\{i \in \mathcal{N}_{t+1}\}} dC_t^{(i)} = C_t. \tag{B.3}$$

This equation follows from the fact that $1_{\{i \notin \mathcal{N}_{t+1}\}} \times 1_{\{i \notin \mathcal{N}_t\}} = 0$ almost surely. That is, the consumption distribution functions $dC_t^{(i)}$ and $dC_{t+1}^{(i)}$ exhibit non-zero increments at different points $i \in [0,1]$. Combining (B.1) - (B.3) along with the goods clearing condition $C_t = Y_t$ we have that

$$\frac{\xi_{t+1}}{\xi_t} = \beta \left(\frac{Y_{t+1}}{Y_t}\right)^{-\gamma} \left(1 - \frac{\int_{i \notin \mathcal{N}_{t+1}} dC_{t+1}^{(i)}}{Y_{t+1}}\right)^{-\gamma}. \tag{B.4}$$

It remains to show that $\frac{\int_{i \notin \mathcal{N}_{t+1}} dC_{t+1}^{(i)}}{Y_{t+1}} = a_{t+1,t+1}$. To show this, we start by noting that the financial wealth of each agent $i \notin \mathcal{N}_{t+1}$ satisfies $\frac{W_t^{(i)}}{W_{t+1}^{(i)}} = 0$. Intuitively, the agent i goes from having infinitesimal wealth to having a point mass of wealth. Therefore, her intertemporal budget constraint gives

$$c_{t+1}^{(i)} \left\{ E_{t+1} \sum_{n=1}^{\infty} \frac{\xi_{t+n}}{\xi_{t+1}} \frac{c_{t+n}^{(i)}}{c_{t+1}^{(i)}} \right\} = a_{t+1,t+1}^{(i)} Y_{t+1} \left\{ E_{t+1} \sum_{n=1}^{\infty} \frac{\xi_{t+n}}{\xi_{t+1}} \frac{Y_{t+n}}{Y_{t+1}} \times \exp\left(-\sum_{k=1}^{n-1} u_{t+k+1}\right) \right\}. \tag{B.5}$$

From this point onwards, the argument follows a "guess and verify" approach. We guess that $\frac{\int_{i \notin \mathcal{N}_{t+1}} dC_{t+1}^{(i)}}{Y_{t+1}} = a_{t+1,t+1}$. Then (B.4) becomes (11). Combining (11) with the first equality in (B.1) implies that $\frac{c_{t+n}^{(i)}}{c_{t+1}^{(i)}} = \frac{Y_{t+n}}{Y_{t+1}} \exp\left(-\sum_{k=1}^{n-1} u_{t+k+1}\right)$ with probability one. Hence (B.5) implies $c_{t+1}^{(i)} = a_{t+1,t+1}^{(i)} Y_{t+1}$, which means that $\frac{\int_{i \notin \mathcal{N}_{t+1}} dC_{t+1}^{(i)}}{Y_{t+1}} = a_{t+1,t+1}$, as conjectured.

Equation (10) follows from simple accounting. Specifically, let $\overline{W}_t^i$ denote the total financial wealth owned by agents $j \leq i$ at time t. Since all agents invest in the same portfolio — the market portfolio — and choose the same consumption-to-wealth ratio, we have

$$\overline{W}_{t+1}^i = \overline{W}_t^i \frac{Y_{t+1}}{Y_t} e^{-u_{t+1}} + L_{t+1}^i \left(1 - e^{-u_{t+1}}\right) \overline{W}_{t+1}^1,$$

which gives

$$\frac{\overline{W}_{t+1}^i}{\overline{W}_{t+1}^1} = \frac{\overline{W}_t^i}{\overline{W}_t^1} \left(\frac{\overline{W}_t^1}{\overline{W}_{t+1}^1} \frac{Y_{t+1}}{Y_t}\right) e^{-u_{t+1}} + L_{t+1}^i \left(1 - e^{-u_{t+1}}\right).$$

This is the desired relation for wealth, given that the term in the large parentheses equals one, and this relation for wealth readily translates to (10), since the weatlh-to-consumption ratio is the same for all agents. ■

Proof of Proposition 2. Let $\mathcal{O}_{t+1}$ be the set of all indices of agents alive at both t and $t+1$ and who do not receive strictly positive-valued firm endowments at $t+1$. For all these agents, the marginal-utility growth is aligned with the pricing kernel, resulting in

$$\frac{\xi_{t+1}}{\xi_t} = \beta \left(\frac{\int_{i \in \mathcal{O}_{t+1}} dC_{t+1}^{(i)}}{\int_{i \in \mathcal{O}_{t+1}} dC_t^{(i)}} \right)^{-\gamma}. \tag{B.6}$$

We further note that

$$\int_{i \in \mathcal{O}_{t+1}} dC_t^{(i)} = (1-\lambda) Y_t \tag{B.7}$$

$$\int_{i \in \mathcal{O}_{t+1}} dC_{t+1}^{(i)} = Y_{t+1} - C_{t+1,t+1} - C_{t+1}^{NR}, \tag{B.8}$$

where $C_{t+1,t+1}$ denotes the time -$t+1$ consumption of the cohort of workers arriving at $t+1$ and C_{t+1}^{NR} is the consumption accruing to the "newly rich" agents at time $t+1$, i.e., the agents who obtained valuable firms. The above three equations imply that

$$\left(\frac{\xi_{t+1}}{\xi_t}\right) = \beta \left(\frac{Y_{t+1}}{(1-\lambda) Y_t} \right)^{-\gamma} \times \left(1 - \frac{C_{t+1,t+1} + C_{t+1}^{NR}}{Y_{t+1}}\right)^{-\gamma}, \tag{B.9}$$

As we already argued in the proof of proposition 1, agents consumption and portfolio decisions are made as if they know they will never be new rich. Consider therefore an agent born at t, who will never become new rich, and who is therefore facing complete markets. The value of this agent's future consumption is

$$E_t \sum_{s=t}^{\infty} (1-\lambda)^{s-t} \frac{\xi_s}{\xi_t} c_{s,t}, \tag{B.10}$$

while his wealth, given by the value of his labor earnings, is

$$E_t \sum_{s=t}^{\infty} (1-\lambda)^{s-t} \frac{\xi_s}{\xi_t} \lambda^{-1} (1-\eta) Y_s \frac{q_{s,t}}{(1-\lambda)^{s-t}}, \tag{B.11}$$

where the term $(1-\lambda)^{t-s} q_{s,t}$ captures the per-capita fraction of aggregate wages accruing at time t to agents born at time s.

Equating these two quantities, we obtain

$$\frac{C_{t,t}}{Y_t} = \lambda \frac{c_{t,t}}{Y_t} = (1-\eta) \frac{\phi_t^l}{\chi_t} q_{t,t}. \tag{B.12}$$

Similarly, treating the new-rich as one representative agent for simplicity, their wealth equals

$$E_t \sum_{s=t}^{\infty} \frac{\xi_s}{\xi_t} \eta Y_s \left(1 - e^{-u_t}\right) e^{-\sum_{t+1}^{s} u_v}, \tag{B.13}$$

which leads to

$$\frac{C_t^{NR}}{Y_t} = \eta \frac{\phi_t^c}{\chi_t} \left(1 - e^{-u_t}\right). \tag{B.14}$$

Equations (B.9), (B.12), and (B.14) give (A.7). ■

Proof of Proposition 3. We derive a steady state of such an economy. In the steady state investors forego $\overline{k}$ units of consumption each period to produce new capital goods. Upon clearing each of the markets for the capital good, we have that

$$Y_t = Y\left(\overline{k}\right) = \overline{A}\left(\overline{k}\right)^{1-\eta},$$

where

$$\overline{A} \equiv \sum_{s \leq t} \left(a_{t,s}\right)^{\eta}.$$

To determine $\overline{k}$ we use two optimality conditions. The first is the familiar Euler equation,

$$\frac{1}{1+r} = \beta E \frac{U'\left(c_{t+1}^{(i)}\right)}{U'\left(c_t^{(i)}\right)} = \beta \frac{c_t}{c_{t+1}}, \tag{B.15}$$

where we have used the assumption of logarithmic preferences and the notation $\frac{c_{t+1}}{c_t}$ to denote the consumption growth of investors that do not receive a valuable firm at time $t+1$. (The reader is referred to the proofs of the previous two propositions for a justification of this Euler equation).

The second optimality condition is that Tobin's q for newly created capital equal one. To formalize this optimality condition, we define the price of capital $q_{t,s}$ as

$$q_{t,s} = r_{t,s}^K + \frac{1}{1+r} q_{t+1,s}, \tag{B.16}$$

where $r_{t,s}^K$ is the time$-t$ rental rate of capital that was produced at time s. In turn, the rental rate of capital is given by its marginal product

$$r_{t,s}^K = (1-\eta)\left(a_{t,s}\right)^{\eta}\left(\overline{k}\right)^{-\eta}. \tag{B.17}$$

Combining (B.16) with (B.17) and $\frac{(a_{t+1,s})^{\eta}}{(a_{t,s})^{\eta}} = e^{-\eta \overline{u}}$, implies a solution whereby $\frac{q_{t,s}}{r_{t,s}^K}$ is constant and therefore $\frac{q_{t+1,s}}{q_{t,s}} = e^{-\eta \overline{u}}$. In that sense $\overline{u}$ is equivalent to economic depreciation of existing capital (reminiscient of the literature on investment-specific shocks). Evaluating (B.16) with $q_{t,t} = 1$ and recognizing that $q_{t+1,t} = e^{-\eta \overline{u}}$ leads to

$$1 = (1-\eta)\, a_{t,t}^{\eta} \left(\overline{k}\right)^{-\eta} + e^{-\eta \overline{u}} \frac{1}{1+r}.$$

Finally, market clearing implies that

$$\frac{c_{t+1}}{c_t} = 1 - \frac{(1-\beta)\,\eta \frac{(a_{t,t})^{\eta}}{\overline{A}} Y \left(1 - \frac{e^{-\eta \overline{u}}}{1+r}\right)^{-1}}{Y - \overline{k}}, \tag{B.18}$$

where we have used the fact that $1-\beta$ is the wealth-to-consumption ratio for a logarithmic investor, along with the fact that the present value of the aggregate rents as a fraction of aggregate consumption is $\frac{\eta \frac{(a_{t,t})^{\eta}}{\overline{A}} Y \left(1-\frac{e^{-\eta\overline{u}}}{1+r}\right)^{-1}}{Y-\overline{k}}$. We note here one difference with the represenative investor setup. Since the rents from new firm creation accrue to a set of measure zero in the population the gross consumption growth rate in (B.18) is not one, but less than one, similar to the baseline model.

Combining (B.18) with (B.15) leads to the following non-linear equation for the determination of $\overline{k}$

$$\frac{e^{-\eta\overline{u}}\beta}{1-(1-\eta)\,(a_{t,t})^{\eta}\left(\overline{k}\right)^{-\eta}} = 1-(1-\beta)\,\frac{\eta}{1-\eta}\frac{1}{\overline{A}\left(\overline{k}\right)^{-\eta}-1}$$

Letting $x = a_{t,t}^{\eta}\left(\overline{k}\right)^{-\eta}$ implies that the above equation is a quadratic equation for x with two solutions in the range $\left[1-e^{-\eta u}, \frac{1}{1-\eta}\right]$. Either of these solutions corresponds to an equilibrium with $\frac{c_{t+1}}{c_t} < 1$, i.e. a consumption growth for almost all agents that lies below the aggregate growth rate. ■

Proof of Proposition 4. Our assumptions imply that the condition for entry into the innovation market is given by

$$(1-\nu)\,M_{t,t} = w_t l_t,$$

where $M_{t,t}$ is the aggregate market value of projects created at time t. Combining the aggregate demand of labor by intermediate firms $l_t^{(I)}$, which can be shown to equatl $\left(\frac{w_t}{\delta^2}\right)^{\frac{1}{\delta-1}}$ (where δ refers to notation introduced in section A.3.1.) with the market clearing condition $l_t + l_t^{(I)} = 1$, and conjecturing that $M_{t,t}$ has the form $M_{t.t} = f(u_t)l_t$, allows one to obtain $w_t = (1-\nu)\,f(u_t)$ and $l_t^{(I)} = \left(\frac{(1-\nu)f(u_t)}{\delta^2}\right)^{\frac{1}{\delta-1}}$. Under this conjecture, and repeating the same steps as in the proof of Proposition 2, the stochastic discount factor is given by

$$\left(\frac{\xi_{t+1}}{\xi_t}\right) = \beta\left(\frac{Y_{t+1}}{Y_t}\right)^{-\gamma} \times \left(1-\frac{\nu}{\chi}M_{t+1,t+1}\right)^{-\gamma}, \tag{B.19}$$

where χ is the consumption-to-wealth ratio, which is conjectured to be constant. Given the conjectures that χ is constant and that $M_{t.t} = f(u_t)l_t$, the stochastic discount factor is i.i.d., which confirms that the consumption-to-wealth ratio is constant, and also that $M_{t,t}$ is equal to $\phi a_{t,t}$ for a constant ϕ, which in turn implies that $\frac{M_{t,t}}{l_t} = \phi\left(1-e^{-u_t}\right)$, consistent with our conjecture. ■

C A calibration of the model with labor income and persistent displacement shocks

In this section of the appendix we provide a model calibration that allows profits and wage income to have different dynamics. Besides showing that our results in section 5 do not depend on this assumption, an additional goal of this section is to illustrate the theoretical implications of including labor income in our model specification. The model that we shall calibrate is the model of section A.2. To illustrate the mechanisms of that model, we start by introducing persistence to displacement shocks – an assumption that is consistent with the data. To obtain closed form solutions, we assume that u_t follows a regime-switching process with two states s_i, $i = [0, 1]$; conditional on the state, the shock u_t is exponentially distributed using different scale parameters, $f(u|s_i) \sim \exp(\theta_i)$. We denote state 0 (1) the low (high) displacement state, that is, $\theta_0 < \theta_1$. We denote the transition probabilities as p_0 (of remaining in state 0 conditional on being in state 0) and p_1 (probability of being in state 1 conditional on being in state 1). These assumptions imply that equations (A.4)-(A.6) become a system of 6 equations in 6 unknowns. We determine the parameters θ_0, θ_1, p_0, and p_1 so as to (approximately) match the magnitudes and shapes of model-implied impulse-response functions in the data (in particular the short- and long- run response of a one standard deviation shock to u_t on index dividends and the divisor) to artificially generated impulse-response functions from the model (the results of this exercise are given in Figure 7). As a robustness check we also estimate θ_0, θ_1, p_0, and p_1 from the time-series on the fraction of market capitalization that is due to new firms, and estimate these 4 parameters by fitting a Markov regime switching model and using the expectation-maximization (EM) algorithm. The two approaches result in essentially identical estimates for θ_0, θ_1, p_0, and p_1, which are given in table 6.

Next, we calibrate the two parameters δ_0 and δ_1 which control the rate of depreciation of human capital in the low- and high-displacement state, respectively. These two parameters, along with the household discount factor β, affect the mean and the volatility of the interest rate. We thus determine δ_0, δ_1, and β to roughly match the mean and the volatility of the interest rate in the data. In addition, we calibrate the mean of δ_i to equal 0.035, effectively treating labor income as a risky bond with a stochastic maturity, which on average equals 33 years.

An alternative interpretation of $\delta_0 > \delta_1$ is that the displacement shocks to labor are drawn from a distribution with a lower mean (compared to the one for firm value) in the high-displacement regime. In that sense, the value of human capital is affected less than the value of firm upon transition to a high-displacement regime. Given the large importance of human capital for the total wealth of the representative investor, such a transition has a muted impact on the expected growth rate of her marginal utility, even though it impacts the volatility and the skewness of her marginal utility. (As already noted, in our model forward-looking valuation ratios affect the moments of the marginal utility of the representative investor). As a result the interest rate remains unaffected by regime transitions, even though the market price of risk is impacted, similar to what would happen in long run risk models with high inter temporal elasticity of substitution.

This ability to control the variation of the interest rate is important both quantitatively, but also conceptually, to illustrate some connections with representative-agent models with recursive preferences. The easiest way to see the issue, is to revisit the simpler model without labor income. If we introduced autocorrelated displacement shocks in such a model, rather than the i.i.d shocks we used, then the presence of CRRA utilities and a risk aversion (inter-temporal elasticity of substitution) larger than one (smaller than one) would have the counterfactual implication of large positive returns in the stock market as the economy shifts to the high displacement regime. This issue is reminiscent of the reason why long run risks models need to assume an inter-temporal elasticity of substitution larger than one.

The remaining parameters are chosen to match standard moments. We choose $\mu = 2.5\%$ and $\lambda = 2\%$, to approximately match the aggregate growth rate of consumption and the population birth rate, respectively. These choices affect predominantly the equilibrium interest rate, and are not important for the pricing of risk. The parameter $\eta = 0.25$ determines the capital share.[20] Last, regarding the parameter governing the volatility of aggregate productivity A_t, we present two versions of the calibration. First, as before, we set $\sigma = 0$, so that all the risk premia are due to re-distributional risks, not aggregate risks. Second, we also calibrate a version of the model with a more realistic assumption of $\sigma = 3\%$. Table 6 reports the parameter values used in our calibration.

The last remaining parameter is risk aversion. We chose a risk aversion of 10 to determine the maximum possible equity premium that can be derived by displacement risk. As we see in Table 7, when $\gamma = 10$, the model can produce an equity premium of around 2% even in a world without any aggregate uncertainty, and an equity premium around 4% if we set $\sigma = 3\%$. In addition, we see that our model can generate not only a plausible equity premium but also a stable and low risk-free rate.

Fluctuations in the volatility of displacement shocks also imply time-variation in risk premia. To gauge the magnitude of this predictability, we estimate predictive regressions of log excess market returns on the log price dividend ratio in simulated data. As we see in Table 8, the model generates a substantial amount of time-series predictability, even though the point estimates are smaller than the data. Nevertheless, there is significant variation in point estimates across simulations, and the empirical values typically lie inside the 95% confidence intervals implied by the model.

In sum, we find that the addition of displacement risk in a simple endowment economy with time-separable preferences helps in terms of explaining not only the equity premium, but

[20]The correspondence between the "capital share" in the data and the "capital share" in the model is not straightforward. Many factors (share of income that is proprietor's income, the presence of real estate, the treatment of depreciation etc.) are present in the data, but not in the model. Moreover, the fact that a substantial fraction of workers are not participating in markets (they are essentially hand-to-mouth consumers) implies that the capital share in the data is understating the fraction of capital income that is directed to active market participants. Fortunately, even though the magnitude of the capital share is sensitive to these assumptions, our results are not. Magnitudes of the capital share between 0.2 and 0.5 lead to similar results – especially if we simultaneously adjust δ_1 and δ_2 to keep the interest rate low and non-volatile.

Figure 7: Impulse Responses: Data vs Model. We simulate 10,000 repetitions of the process for dividends and the divisor in equations (6) and (8), where the shock u_t is exponential with two different scale parameters θ_1, θ_2 that change in a regime-switching fashion. Using the parameters of table 6 and setting the volatility of the neutral component σ to 0.07 (so as to produce realistic implications for the impulse response functions of the neutral shock on the two endogenous quantities) we draw 10,000 artificial paths of length equal to the data. The dashed lines report the 95% bands and the mean of the impulse responses in these simulations. The solid line corresponds to the data (see Figure 4).

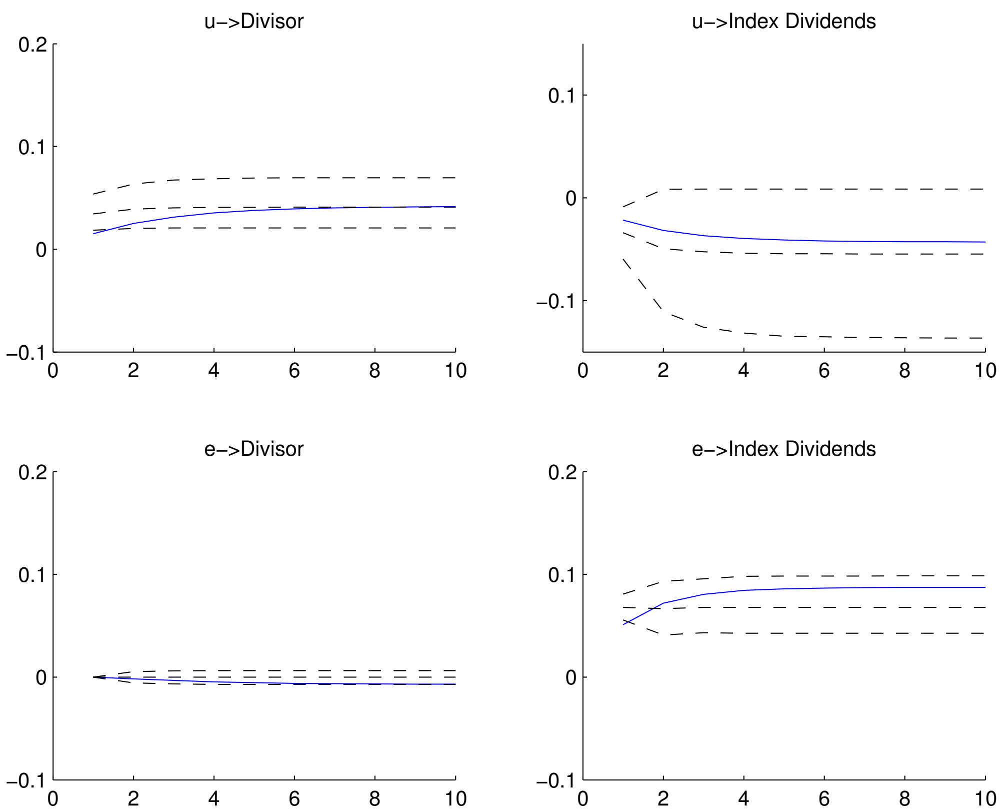

Table 6: Parameters used in calibration of extended model. We report the parameters for the model described in Section A.2.

β	0.970	γ	10	λ	0.020
μ	0.025	σ	0.000	p_0	0.850
θ_0	0.01	θ_1	0.045	p_1	0.85
η	0.250	δ_0	0.070	δ_1	0.000
ψ	1				

also the risk free-rate puzzle, while also being consistent with the stylized facts of non-volatile real rates and return predictability.

Table 7: Comparison between the model and the data. The last column gives results when we set $\sigma = 0.03$. Since $\sigma = 0.03$ introduces a motive for precautionary savings, to keep the models comparable we adjust the subjective discount factor to keep the average interest rate roughly unchanged compared to the baseline model. For the results on the equity premium we multiply the un-levered equity premium by 1.6, consistent with historical leverage ratios. The data column is taken from Gârleanu and Panageas (2015). We refer to that paper for details on the data sources.

	Data	Baseline model	Model ($\sigma = 0.03$)
Average Interest rate	0.028	0.031	0.025
Average Price-dividend ratio	26	22.06	19.1
Average Equity premium	0.052	0.021	0.040
Average Sharpe Ratio	0.29	0.306	0.503
Volatility of Interest rate	0.009	0.008	0.014

Table 8: Long-horizon regressions of excess returns on the log P/D ratio. The simulated data are based on 1000 independent simulations of 100-year long samples, where the initial state is drawn from the stationary distribution of high- and low-displacement states. For each of these 100-year long simulated samples, we run predictive regressions of the form $\log R^e_{t+h} = \alpha + \beta \log \frac{P_t}{D_t}$, where R^e_{t+h} denotes the time-t gross excess return over the next h years. We report the mean values of the coefficient β and the R^2 in these simulations along with [0.025, 0.975] percentiles.

Horizon(Years)	Data (β)	Model (β)	Data (R^2)	Model (R^2)
1	-0.13	-0.114	0.04	0.015
		[-0.469 0.255]		[0.000 0.070]
3	-0.35	-0.243	0.090	0.029
		[-1.135 0.731]		[0.000 0.128]
5	-0.60	-0.291	0.18	0.036
		[-1.539 1.157]		[0.000 0.151]
7	-0.75	-0.336	0.23	0.040
		[-1.858 1.325]		[0.000 0.174]

D The implications of a misspecified orthogonality condition

We conduct the following thought experiment: Suppose that instead of using our decomposition approach, we simply postulated one element of the matrix B. Since B captures the immediate impact of the structural shocks on the system, we choose that element of B to be the immediate response of log dividends per share to the displacement shock.

Conditional on this choice, we can identify the rest of the matrix B without utilizing our instrument. Having obtained B we can determine the dynamics of the entire system. Taking B to be the truth, suppose we are given a problematic instrument $z_t = u_t + \phi\epsilon_t^{p\perp}$ and use our procedure along with this problematic instrument to infer $\hat{B}(B, \phi)$. Letting $\hat{B}^{\text{data}}$ denote the estimate of B that we obtain in the data, we choose ϕ so that

$$\hat{B}^{\text{data}} = \hat{B}(B, \phi)$$

In other words, we set ϕ so that the estimated matrix B would be exactly the same as the one we obtain from the data, even though the true B is different. Using this estimated B, we apply our decomposition method to extract the two permanent shocks from the series and obtain the correlation between the inferred displacement shock and the two actual shocks.

We repeat the above procedure for many different postulates on what is the true value of the response of log dividends per share to the displacement shock. We choose the range of these postulates starting from the assumption that our instrument is correct, and then considering what would happen as the true response of log dividends per share to the displacement shock becomes weaker.

The figure plots the correlation between the inferred and the actual displacement shock, and the correlation between the inferred displacement shock and the actual permanent neutral shock.

The figure shows that if our instrument is problematic, then the inferred displacement shock will be *positively* correlated with the actual permanent neutral shock.This observation is important for our purposes: In the next section we find that our inferred displacement shock has a negative price of risk (high u signals "bad times"). Assuming that $\epsilon_t^{p\perp}$ has a positive price of risk (which would hold for any increasing, concave utility function) then Figure 8 implies that a problematic instrument would make us understate the true market price of displacement risk, since our inferred displacement shock would commingle a (positively priced) neutral shock.

In results not reported here, we also found that the fraction of the long run variance of dividends per share that is due to the permanent neutral shock would have to be extremely low (around 0.2) if the correlation between actual and inferred displacement shock fell below 0.85. So unless someone is willing to accept that the displacement shock explains the bulk of the long-run variation of dividends (a conclusion that even we would find implausible), it seems unlikely that our instrument is problematic.

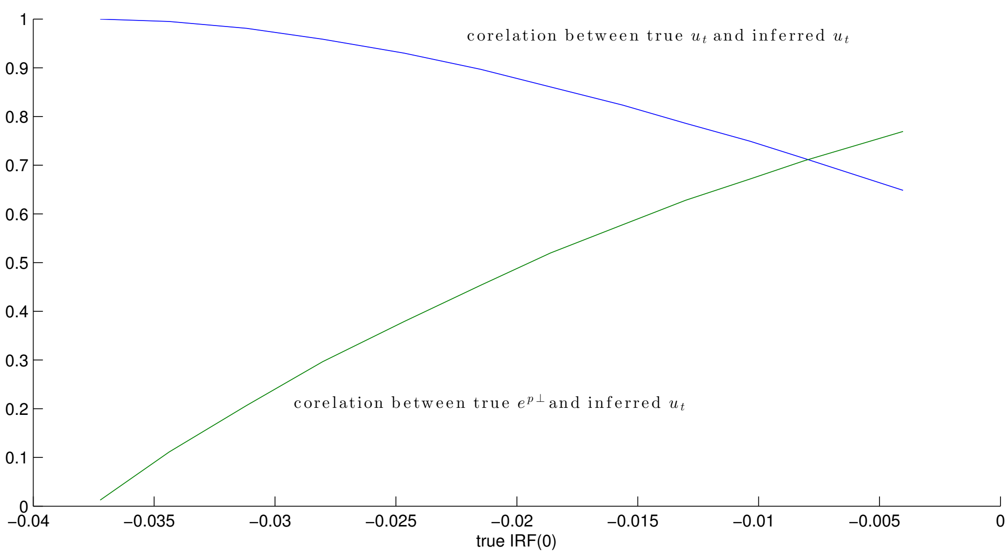

Figure 8: Correlation between true and inferred displacement shock (u_t) and correlation between inferred displacement shock and true $\epsilon_t^{p\perp}$ for various assumptions on the true instantaneous impact of the displacement shock on log dividends per share (IRF(0)).

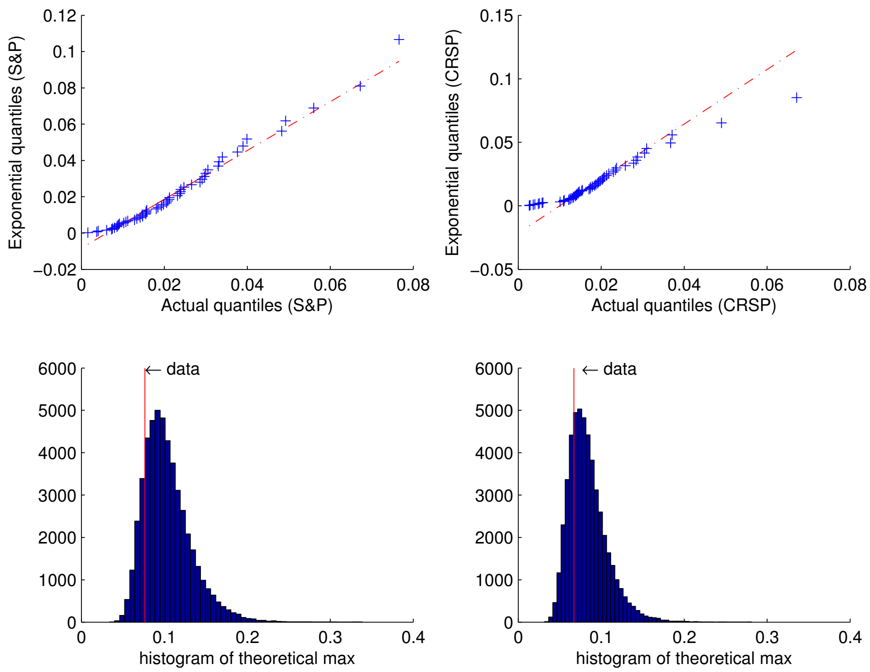

Figure 9: The top left (right) plot is a q-q plot of quantiles of the fraction of new company market value as a fraction of aggregate market valuation for the S&P 500 (left) and the entire CRSP universe (right) dropping the years when AMEX and NASDAQ enter the sample. The bottom figures report results of a Monte Carlo exercise. We draw 56 values (the length of our data) from an exponential distribution, with scale parameter estimated via maximum likelihood (S&P data on the left, CRSP on the right). We record the maximal value, repeat the exercise 10,000 times, and then compare the distribution of the maximal values to the respective maximal value in the data (vertical line labeled "data"). The line "data" is well above the 10-th percentile of the Monte Carlo simulations and indeed close to the mode.