Investment Flexibility and Stock Returns

Dirk Hackbarth and Timothy Johnson*

November 2011

Abstract

Firms with greater investment flexibility do not generally have safer stock or lower expected returns. In a model of firm expansion and contraction, we show that a distinguishing feature of firms with lower adjustment costs (i.e., higher flexibility) is that their risk falls on average as profitability declines and operating leverage (or book-to-market) increases, whereas risk rises with operating leverage for firms with less flexibility. Empirical tests provide support for the model's predictions. The evidence is consistent with cross-firm heterogeneity in investment flexibility, with most firms having valuable contraction options. This makes it unlikely that irreversibility explains the value premium.

KEYWORDS: REAL OPTIONS, FLEXIBILITY, RISK PREMIA, VALUE EFFECT.

JEL CLASSIFICATIONS: D31, D92, G12, G31

---

*Both authors are from the University of Illinois, Urbana-Champaign. We are very grateful to Jonathan Berk, Lorenzo Garlappi, Graeme Guthrie, Ali Ozdagli, Tano Santos, and Lu Zhang for their thoughtful comments. We also thank seminar participants at the 2011 NBER Asset Pricing Meeting, the University of Calgary, the University of Georgia, the University of Illinois, and the University of Rochester. Lifeng Gu provided excellent research assistance.

1 Introduction

Do more valuable real options make stock returns safer and thereby lower expected returns? More generally, how do the risk properties of firms' equity differ for differing degrees of investment flexibility? While the real options literature has long recognized that variation in investment and disinvestment costs can imply important differences in optimal investment policy (see, e.g., Abel, Dixit, Eberly, and Pindyck (1996) and Abel and Eberly (1996)), the implications of this heterogeniety has received little attention in the asset pricing literature.

Research to date has focused on models of *ex ante* homogeneous firms that differ only in their history of idiosyncratic shocks. Homogeneity is a useful simplifying assumption in otherwise complex models, and also enables the isolation of effects that come solely through productivity differences. This line of research has produced numerous insights and delivered several successes in explaining particular observed patterns in stock returns.[1] To the extent that the models derive their results from variation in risk that stem from changes in the relative value of firms' real options, it is natural to ask what testable implications arise from differences in option values across firms (as well as over time). The empirical investment literature has documented substantial differences across firms in the purchase and resale costs of physical capital.[2] These differences are equivalent to differences in the unconditional value of expansion and contraction options. This study is among the first to explore the effect of cross-firm differences in real option values for the risk and return characteristics of equity.

We employ a model that is both rich enough to encompass wide variability in firm parameters, and yet simple enough to reveal general implications. The model is set in partial equilibrium with a constant riskless rate and market price of risk. The single state-variable is the firm's instantaneous operating profit scaled by its capital, which

---

[1] See Carlson, Fisher, and Giammarino (2004), Zhang (2005), Cooper (2006), and Li, Livdan, and Zhang (2009) among others.

[2] See for example MacKay (2003), Balasubramanian and Sivadasan (2009), and Chirinko and Schaller (2009).

is monotonically related to Tobin's Q, or inversely related to the book-to-market ratio, B/M.

Interestingly, almost irrespective of parameter values, the graph of expected excess returns as a function of B/M describes a characteristic S-shape. The slope of the mid-section of this graph may have either sign, but the end sections almost always display an increasingly negative slope. This pattern is due to the firm's real options. Contractions options decrease overall risk as B/M rises, while expansion options increase overall risk as B/M declines.

Our first conclusion is that, perhaps contrary to intuition, the unconditional relationship between investment flexibility and equity risk premia need not be negative. Lowering some types of adjustment cost can raise average expected returns. The finding is not obvious: computing the unconditional effects requires not only solving the firms' problem, but integrating over the (endogenous) distribution of operating states.

While the *level* of the risk premium is not, in general, increasing in measures of inflexibility, the *slope* of the risk premium is. Our second finding is that the sign of the relationship between profitability (or Q) and expected returns may go either way, depending on the adjustment cost parameters. Again perhaps surprisingly, more flexible firms actually become safer as profitability declines. This is because their option to exchange risky physical assets for safe cash becomes increasingly valuable as productivity deteriorates. While, in principle, the model is not inconsistent with the assertion that the book-to-market effect in the cross-section of returns is primarily driven by irreversible investment, the logic breaks down with even moderate degrees of down-side flexibility.

In interpreting the model, we note that the notion of flexibility extends beyond physical capital. One can interpret the firm's production function as dependent on a general factor input which can be viewed as bundle of capital, labor, and firm-specific knowledge. Further, the quasi-fixed costs in the model are not just associated with physical capital but may accrue from any inputs purchased with long-duration contracts.

Turning to the data, these observations suggest that firm scale flexibility should not be assessed purely with respect to investment activity. Instead, viewing real technologies as largely industry-specific, we construct four industry-level, measures of inflexibility from estimates of overall cost stickiness and variability. These measures are either directly derived from or closely linked to the model. Similarly, we assess a given firm's period-specific operating leverage (or profitability) based on expected quasi-fixed costs over sales, without regard to the source of these costs.

With these measures, we verify the primary cross-sectional predictions of the model. Consistent with the model, we find no unconditional flexibility effect in the data. Portfolios formed on sorts of industry flexibility do not support the conjecture that more flexible firms have unambiguously lower returns. Sorting within industries by operating leverage, however, does reveals the predicted difference in slope between more and less flexible sectors. In a cross-sectional regression framework, these finding are robust to the inclusion of standard controls and to alternative measurement of both the conditioning variables. Notably, the interaction effect is consistent with significant cross-firm heterogeneity in real options. The sign of the marginal flexibility effect is negative, suggesting that on average firms have valuable contraction options.

To summarize, we demonstrate that differences across firms in investment flexibility do lead to economically significant differences in the risk and return characteristics of their equity. Crucially, however, the pattern predicted and (verified in the data) is not about the unconditional level of expected returns but about the response of the risk premium as profitability (or operating leverage) evolves. Interestingly, neither the operating leverage proxies, the flexibility proxies nor their interaction lowers the explanatory power of the book-to-market ratio[3], suggesting that the value effect is more likely to be driven by cross-firm differences in asset risk than by within-firm variation in operating leverage.

The outline of the paper is as follows. The next section introduces the model and derives its implications. Section 3 describes how we operationalize these predictions empirically. Section 4 presents our empirical results, and Section 5 concludes.

---

[3] This is also the case in the results of Garca-Feijó and Jorgensen (2010) and Novy-Marx (2011), both of whom find some explatory power for measures of operating leverage in the cross-section.

2 The Model

We employ a neoclassical investment model in partial equilibrium to investigate the expected return implications of scale flexibility. A number of related models have been used in the literature; ours has appealing generality and tractability properties that make it well-suited for studying complex questions. We use it to analyze how the firm's real technology maps into (a) its expected stock return as a function of its state, and (b) its endogenous distribution of operating states.

The setting is a continuous-time economy with a fixed interest rate r and a pricing kernel, Λ, that obeys a geometric Brownian motion with volatility σ_Λ that characterizes the economy's risk-reward trade-off. Each firm in the economy is a claim to a real production function characterized by declining returns to scale and quasi-fixed operating costs.

The scale of the firm is denoted K. As described in the introduction, we prefer to think of K as a bundle of productive factors that the firm has in place, encompassing anything that generates quasi-fixed operating costs, such as, labor inputs or organizational capital. These costs, which are proportional to K but do not scale with output, are the main mechanism generating return dynamics. So the economic logic of the model is not confined to physical capital. Without loss of generality, though, physical capital is taken to be the numeraire. So K can be viewed as the book value of assets.

At each point in time, the firm's output – or revenues net of variable costs – are determined by K together with the level of productivity θ. The productivity process evolves according to a geometric Brownian motion with drift μ, volatility σ and correlation with the pricing kernel denoted ρ. The firm's profit flow (per unit time) is

$$\Pi_t = \theta_t^{1-\gamma} \; K_t^{\gamma} - m \, K_t, \tag{1}$$

where $\gamma \in (0,1)$ captures returns to scale of the firm and $m > 0$ denotes the firm's operating cost per unit of K. Unless adjusted by the firm, K follows $dK/K = -\,\delta\, dt$, with the depreciation rate $\delta \geq 0$. The model is partial equilibrium both because the kernel is exogenous and because we do not model interactions between firms.[4] Note also that for tractability we consider only permanent productivity shocks.

For present purposes we confine attention to an all equity-financed firm. Recently Ozdagli (2011) has analyzed a version of the model studied here for a firm with debt. In a setting in which it is costly for the firm to deviate from a constant book leverage, interests costs act to magnify the quasi-fixed operating costs. In an appendix available upon request, we verify that the primary features we describe here for *firm* expected returns are preserved for *equity* expected returns under some reasonable formulations of debt determination.

Operational flexibility has several components for real firms, including the ability to change pricing strategies, enter new businesses, or alter production technologies. In the present model, *investment* flexibility consist of the ability to increase or decrease scale in response to shocks to profitability. We assume firms face both quasi-fixed and variable costs for either upward or downward adjustments.

When increasing the scale of operations, the firm faces costs that are proportional to net revenue at the time of the adjustment, $F_L\ \theta^{1-\gamma}\ K^{\gamma}$, where $F_L \geq 0$ (the subscript will be explained below). In addition, the cost to investors of increasing K by ΔK may exceed ΔK, e.g., due to installation frictions. These costs are assumed linear: the amount required from investors is $P_L\ \Delta K$ where $P_L \geq 1$. The deadweight loss from the adjustment is thus $(P_L - 1)\ \Delta K$.

We assume that the frictions for disinvestment are the same as those for investment. Specifically, for any contraction of scale there are fixed costs denoted $F_U\ \theta^{1-\gamma}\ K^{\gamma}$. And the cash returned to investors when K is lowered by ΔK is taken to be $P_U\ \Delta K$, with $P_U \leq 1$. In principle, we could even have $P_U < 0$ due, e.g., to penalties for breaking contracts with suppliers. Note that our assumptions here do not actually nest the case of irreversibility. The option to disinvest – even with little payoff – is still better than

---

[4] General equilibrium models of investement-based return effects appear in Gomes, Kogan, and Zhang (2003), Gala (2006) and Sagi, Spiegel, and Watanabe (2009). Industry competition is considered in Zhang (2005), Aguerrevere (2009), and Novy-Marx (2011).

no option except abandonment.[5] To our knowledge, this is the first real-options model to incorporate the ability to repeatedly expand and contract under this cost structure.[6]

Because of the frictions, the firm pursues a discrete adjustment policy. Specifically, with a given level of K, it will increase to $K' > K$ only when productivity attains some level $\theta_L(K)$. But, since the profit function and adjustment costs are all homogeneous of degree one in assets, once at K' the firm faces an identical environment scaled up by the ratio K'/K. It follows that both θ_L and K' are proportional to K. By a similar argument, disinvestment will occur only when θ falls to some θ_U proportional to K and the disinvestment will lower assets to some K'' also a fixed fraction of the prior K. The firm's problem is to choose the four ratios $\theta_L/K, K'/K, \theta_U/K, K''/K$ to maximize the expected discounted sum of future profits under the risk-neutral measure. Equivalently, following Cooper (2006), if we define $Z_t \equiv K_t/\theta_t$, then the four constants correspond to four points on the Z axis: investment happens at the lower boundary $Z_t = L = K/\theta_L$ and moves the firm to $Z_t = G = (K'/K)L > L$; disinvestment happens at the upper point $Z_t = U = K/\theta_U$ and moves the firm to $Z_t = H = (K''/K)U < U$. The firm thus lives on the interval $[L, U]$. In terms of the original variables, the firm's path in the $K-\theta$ plane, depicted in Figure 1, describes oscillations along lines of fixed K between two rays $K = U\theta$ and $K = L\theta$ with jumps up and down to the interior rays $K = G\theta$ and $K = H\theta$. (The figure sets the depreciation rate to zero for simplicity.)

The effect of adjustment costs on the optimal policy is straightforward. A frictionless firm with no adjustment costs will, given θ, set K to the value $(m/\gamma)^{1/(\gamma-1)}\ \theta$ that maximizes the profit function Π. Denote the K/θ ratio at this point $Z^{\star}$. Now as fixed or variable investment costs are increased, the firm will choose a smaller value of $L < Z^{\star}$, waiting longer between adjustments. Likewise, either type of cost for disinvestment raises $U > L$.[7] In terms of the model, a good summary statistic for scale inflexibility is the

---

[5] We do allow for abandonment, however. For each firm, we solve the valuation problem with costly disinvestment and if it entails negative firm value we re-solve imposing the boundary conditions for abandonment – which ensures nonnegativity – instead.

[6] Cooper (2006) studies the case of purely irreversible investment. Guthrie (2010) incorporates a one-time disinvestment option in a similar setting.

[7] Intuition might suggest that raising the variable costs of adjustment leads to smaller changes, mean-

Figure 1: *Firm Evolution*

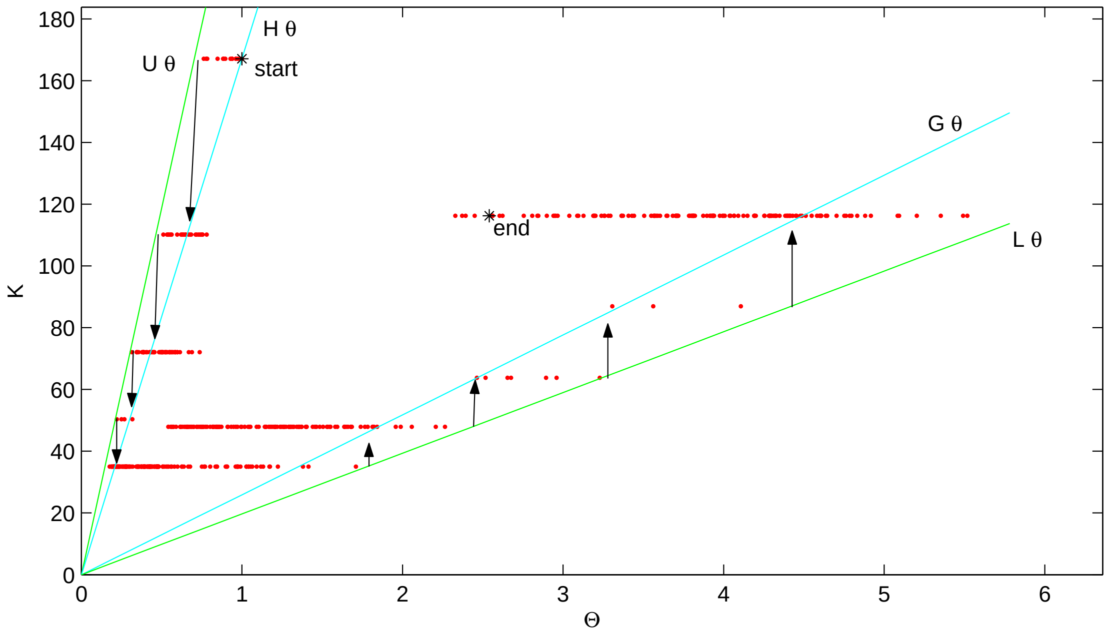

The figure shows a simulated path of a model firm in the K-θ plane. The firm parameters are $\gamma = 0.85, m = 0.4, \delta = 0.0, P_L = 1.0, F_L = 0.01, P_U = 0.25, F_U = 0.01, \mu = 0.05, \sigma = 0.3, \rho = -0.5$. The pricing kernel has $r = 0.04$ and $\sigma_\Lambda = 0.50$.

distance between the adjustment boundaries, $\log(U/L)$, standardized by the volatility of the productivity process σ.

The value of the firm at any time can be written $J(K, \theta) = \theta\, V(Z)$ where V is given in closed-form in the appendix. Fully characterizing the dependence of this function the firm parameters is not possible, because it depends on the solution of a six-equation algebraic system that is not expressible analytically in terms of the production and adjustment cost variables. Two intuitive properties of the solution are the following:

(A) The market-to-book ratio rises monotonically with θ, and hence V/Z falls with Z.

(B) If abandonment is never optimal, the ability to adjust operations buffers firm risk,

---

ing G and H close to L and U respectively. However this is not the usually the case: the desire to avoid paying the costs more frequently counteracts the incentive to minimize the adjustment. Another intuition that is not supported is that higher costs lead to longer periods between adjustment. The first-order determinants of how frequently an adjustment barrier is hit are the stochastic parameters, μ and σ and the depreciation rate, δ.

specifically $(\theta/J)(\partial J/\partial\theta) < 1$.

The expected excess return to the firm's equity (or risk premium) is given by

$$EER(Z) = \pi_\theta\,(1 - Z\,V/V')\,, \tag{2}$$

i.e., the elasticity of J w.r.t. θ times $\pi_\theta \equiv -\,\rho\,\sigma\,\sigma_\Lambda$, the market price of θ-risk. Assuming $\pi_\theta > 0$, Property (B) implies that $EER(Z) < \pi_\theta$, and Property (A) is equivalent to $EER(Z) > 0$.

Given the lack of analytical characterizations of the effects of firm parameters on expected returns, we illustrate the properties we identify throughout the paper by solving a large number of cases. Specifically, the model is solved with each of the 2^9 combinations of parameters shown in Table 1, which covers a large range of firm characteristics while staying with in the bounds of plausible expected returns and volatility.[8] In none of these 512 cases is abandonment optimal. In all of the cases Properties (A) and (B) are satisfied.

Table 1: *Parameter Ranges*

	PRODUCTION:		FRICTIONS:				STOCHASTIC:		
Range:	γ	δ	P_L	F_L	P_U	F_U	μ	σ	ρ
High value	0.75	0.00	1.50	0.05	0.60	0.05	0.04	0.55	-0.10
Low value	0.95	0.10	1.00	0.005	0.10	0.005	0.00	0.25	-0.90

The table shows the range of parameters considered in numerical verification of the assertions in the text. Each of the 2^9 combinations of high and low values are computed. Each case sets the parameter m to $\gamma\ 100^{\gamma-1}$ which puts $Z^\star = 100$. All cases use $r = 0.04$ and $\sigma_\Lambda = 0.50$.

As we show in the appendix, the risk premium, $EER(Z)$, can be decomposed into three distinct components, namely, the exposures due to assets in place and the operating options to expand or contract. Figure 2 charts these three components for a typical

---

[8] There are actually 10 firm-specific parameters in the model. However, given the returns-to-scale parameter γ, the cost parameter m acts mainly to scale the problem on the Z axis. (The units of Z are otherwise arbitrary.) Hence we fix $m = \gamma\ 100^{\gamma-1}$ which puts $Z^\star = 100$ for all cases.

case. The horizontal axis, instead of Z, is $mZ^{1-\gamma}$, which is the ratio of quasi-fixed costs, mK to net sales (or operating margin) $\theta^{1-\gamma}K^{\gamma}$. This quantity, which we denote QFC/S, is essentially a measure of the degree of operating leverage at any point in time. It is convenient because it is positive, monotonic in Z, and plausibly measurable in the data. The figure reveals that, while the risk from assets in place monotonically increases with quasi-fixed costs (i.e., increases with operating leverage), the risk from *both* operating options declines with quasi-fixed costs (i.e., rises with profitability). Note that the disinvestment option attenuates exposure to priced risk and the investment option exacerbates exposure to priced risk. The contrasting effect of these options is the key feature of the model.

Figure 2: *Components of Risk Premium*

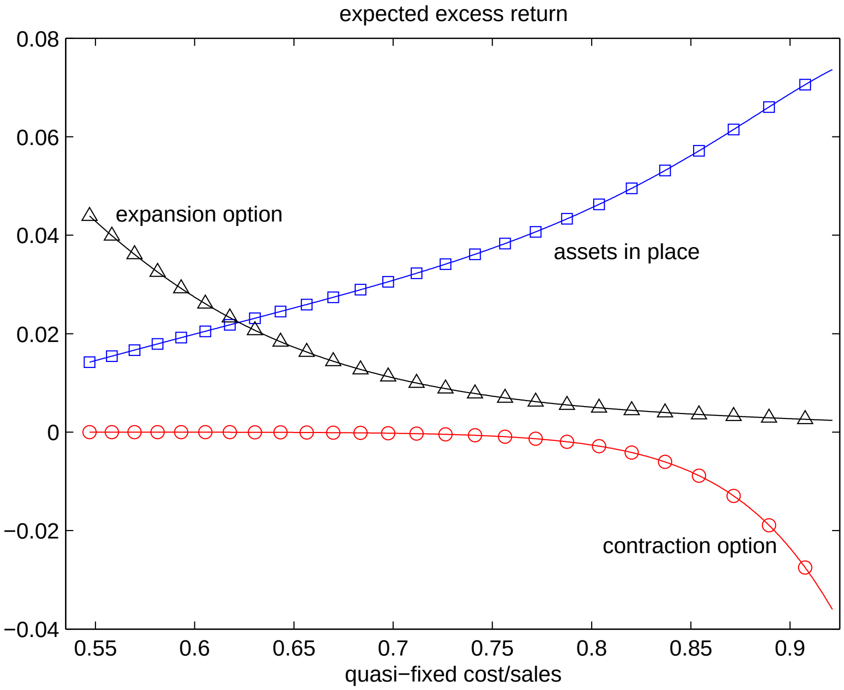

The figure shows the three components of expected excess returns as a function of the ratio of quasi-fixed costs to net sales for a firm with assets in place (plotted as squares), contraction option (circles), and expansion option (triangles). Firm parameters are $\gamma = 0.85, m = 0.4, \delta = 0.1, P_L = 1.0, P_U = 0.25, F_L = 0.05, F_U = 0.05, \mu = 0.05, \sigma = 0.3, \rho = -0.5$. The pricing kernel has $r = 0.04$ and $\sigma_\Lambda = 0.50$.

We show in the appendix that the signs of the slopes of the three risk components are general properties that are satisfied for all parameter values for which (A) and (B) hold.

In sum, the three components imply a distinctive sideways S-shaped plot of expected return versus Z (or QFC/S). In all the solutions of all the cases of Table 1 the plots exhibit negative slopes at L and U, with a switch to a more positive slope in the middle.[9]

We now examine the effect investment flexibility (i.e., adjustment costs) on the risk-return profile. It turns out that the effects of the variable-costs parameters predominate.[10]

Figure 3: *Effect of Reversibility*

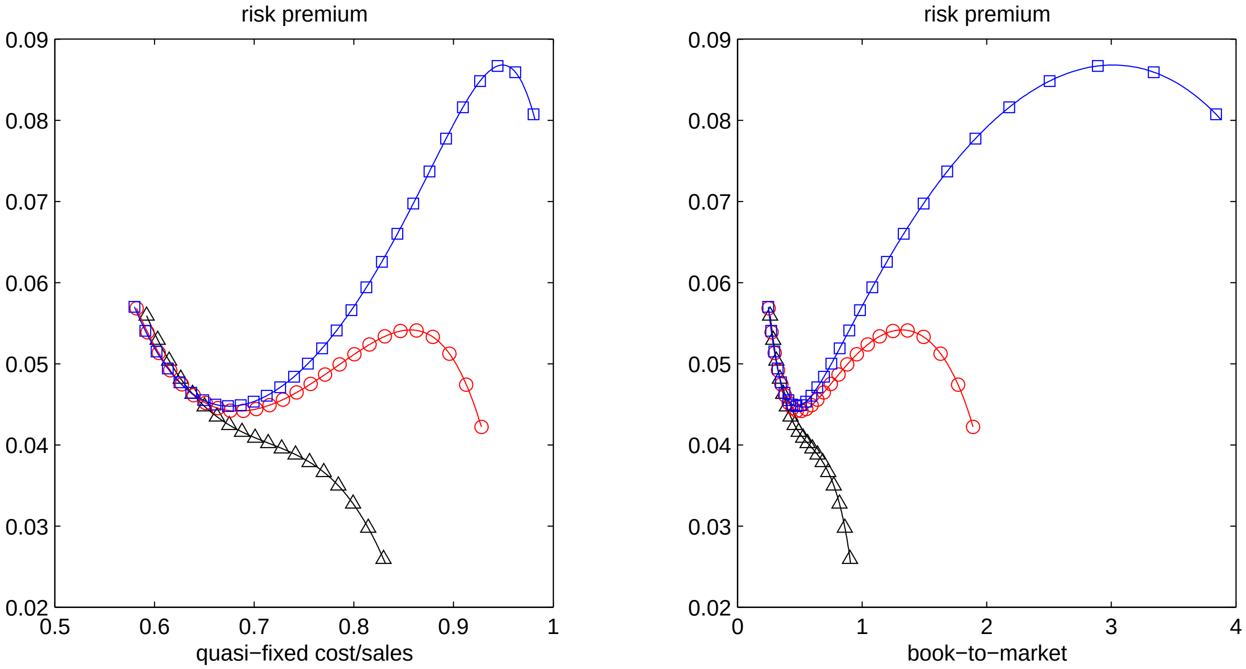

The figure shows expected excess returns for firms with $P_U = 0.01$ (plotted as squares), $P_U = 0.25$ (circles), and $P_U = 0.6$ (triangles). In the left panel the horizontal axis is the ratio of quasi-fixed cost to net sales; in the right panel it is the book-to-market ratio. Other firm parameters are $\gamma = 0.85, m = 0.4, \delta = 0.1, P_L = 1.0, F_L = 0.05, F_U = 0.05, \mu = 0.05, \sigma = 0.3, \rho = -0.5$. The pricing kernel has $r = 0.04$ and $\sigma_\Lambda = 0.50$.

Consider Figure 3, which shows the effect of changing the degree of reversibility, P_U, on expected excess returns for a particular case (the parameters are given in the figure caption). The left panel again uses QFC/S on the horizontal axis; the right panel uses

---

[9] In 70 percent of the cases, the slope in the middle is positive, implying the distinctive S-shape with one inflection point. For the other 30 percent, the slopes are everywhere negative, but still feature the three well-defined regions. If abandonment did dominate contraction, then the S-shape would turn into a J-shape because the disinvestment option component would vanish.

[10] The fixed cost parameters, F_U and F_L play much less significant roles in determining risk profiles. For plausible parameter ranges, (e.g. under 5 percent of net sales) the total contribution of these costs is relatively small and they are incurred infrequently.

the book-to-market ratio, $K/J = Z/V$. As both panels illustrate, making the firm's technology more irreversible by lowering P_U has two effects. First, overall it makes the stock riskier and raises the expected excess return. Second, it raises the average *slope* of the curve: expected excess returns rise steeply with operating leverage (at least over the middle part of the graph) for firms with nearly irreversible capital stocks. Both effects stem from shrinking the contribution of the disinvestment option component of the risk premium as P_U is lowered. Note that, as P_U declines (irreversibility increases), the right end-point of the plots shift out: the firm endogenously chooses to increase U – putting its disinvestment option further out of the money.

The expected return pattern for low P_U firms are consistent with existing findings in the literature based on irreversible investment. What is novel, however, is that for firms with even a mild degree of reversibility, the average slope of the line is negative: the stock actually becomes safer as profits decline and operating leverage increases. For these firms, the contribution of the disinvestment option actually overwhelms the effect of operating leverage.[11] Note that the model implies that within-firm variation in profitability or operating leverage should imply an *anti*-value effect for these firms.

The discussion shows the crucial role that the liquidation cost parameter, P_U, plays in determining the relative contribution of the disinvestment option to firm risk. Likewise, the key parameter determining the strength of the expansion option is P_L, the effective cost of a unit of capacity. Figure 4, shows a typical case. The first message is that the slope conclusion continues to apply: it is still true that a less flexible firm exhibits a steeper (or more positive) average increase in risk premium with operating leverage. However, here it is *not* the case that the plot for the higher P_L firm is everywhere higher than for the lower P_L firm. In this case, inflexibility is not unconditionally associated with equity risk. This finding runs somewhat counter to intuition: the firm does not utilize its

---

[11] In a similar model, Guthrie (2010) analytically shows the negative dependence of expected returns on operating leverage for the case of a firm with a one-time abandonment option, but otherwise fixed scale. The intuition in his case is identical to that in our model. Moreover, the idea is related to the effect in Garlappi, Shu, and Yan (2008) and Garlappi and Yan (2011) where firms approaching bankruptcy experience decreasing risk premia if the the absolute priority rule is violated and hence equity holders can extract (less risky) recoveries instead of nothing.

Figure 4: *Effect of Investment Costs*

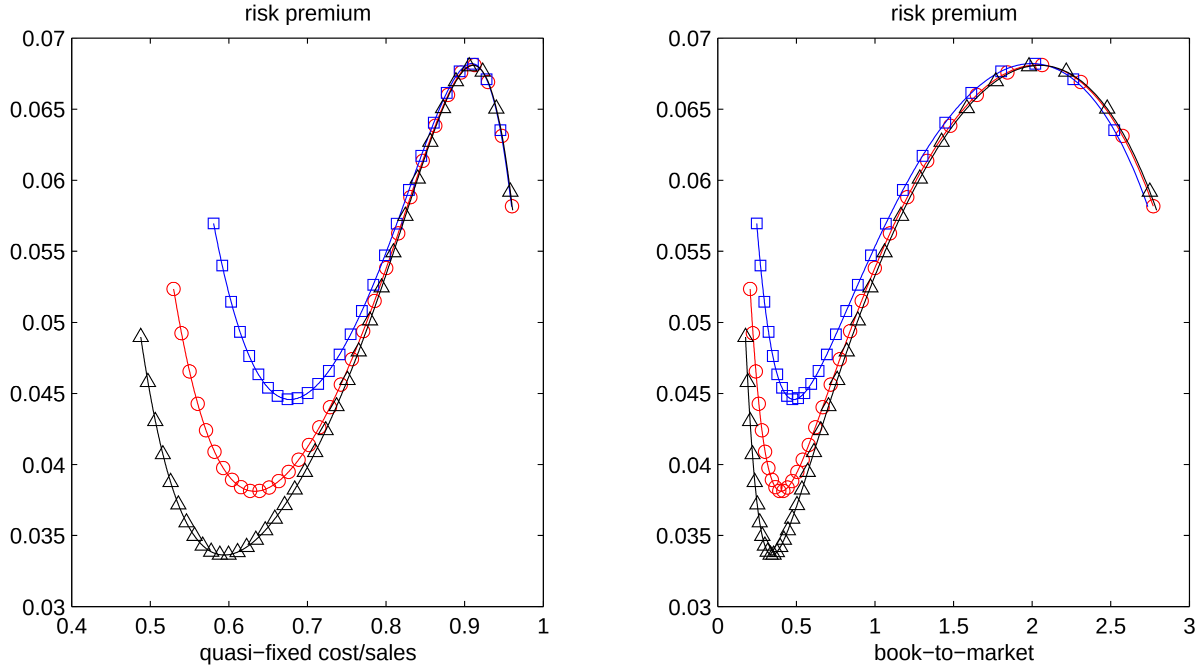

The figure shows expected excess returns for firms with $P_L = 1.0$ (squares), $P_L = 1.5$ (circles), and $P_L = 2.0$ (triangles). In the left panel the horizontal axis is the ratio of quasi-fixed cost to net sales; in the right panel it is the book-to-market ratio. Other firm parameters are $\gamma = 0.85, m = 0.4, \delta = 0.1, P_U = 0.25, F_L = 0.05, F_U = 0.05, \mu = 0.05, \sigma = 0.3, \rho = -0.5$. The pricing kernel has $r = 0.04$ and $\sigma_\Lambda = 0.50$.

growth options to buffer the exposure of its dividends to exogenous profitability shocks.

Characterizing the general effect of scale flexibility on expected returns requires taking into account not just the plot of the function but also the region of the horizontal axis in which the firm operates. For example, a firm with valuable contraction options (hence lower overall risk) may also have parameters leading it to tend towards high-Z states where its risk is high. We thus simulate time series for our 512 model cases and compute unconditional expected equity returns for each.

Figure 5 illustrates our finding that, in general, firm flexibility is *not* a determinant of the average level of expected returns. The plot shows the unconditional risk premium against the summary measure of flexibility $\sigma^{-1}\log(U/L)$, which incorporates the contribution of both expansion and contraction options. The plot uses different symbols to distinguish cases of different systematic *asset* risk, measured by the product $-\rho\,\sigma$.

Whether one controls for these differences or not, there is clearly not a positive relationship between the variables plotted.

Figure 5: *Flexibility and Unconditional Expected Returns*

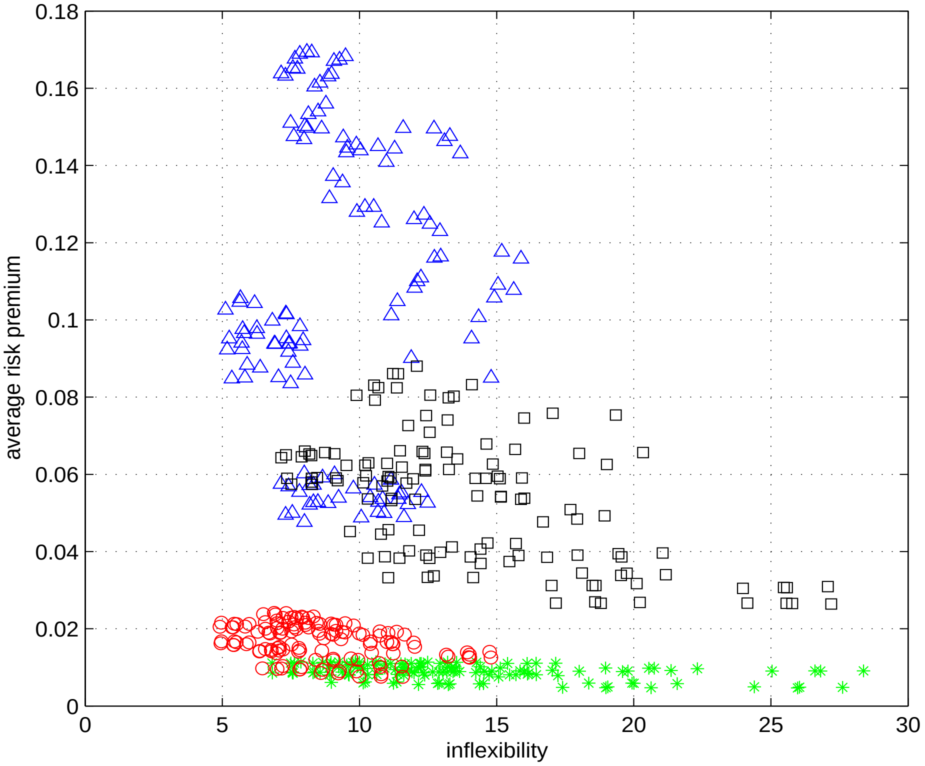

For each of the models described in Table 1, unconditional expected excess returns are computed by integrating the theoretical risk premium with respect to the distribution of the state variable Z obtained from 1000 year simulations. The result is plotted against a summary measure of the firm's flexibility (the scaled range of its no-adjustment region). Cases with $\sigma = 0.55, \rho = -0.9$ are plotted as triangles. Cases with $\sigma = 0.25, \rho = -0.9$ are plotted as circles. Cases with $\sigma = 0.55, \rho = -0.1$ are plotted as squares. Cases with $\sigma = 0.25, \rho = -0.1$ are plotted as asterisks.

To be clear, the theory is not inconsistent with there being a return premium in the data for inflexible firms. But if there is, the correct inference is not that this is an intrinsic property of the neoclassical model, but rather that the parameter heterogeneity in the data happens to include firms that tend to plot on a diagonal in Figure 5.

Figure 6 uses the same simulations to compute the average *slope* of the expected return graph for each firm. For each firm's history, the slope is determined from the regression of true expected returns (sampled daily) on its operating leverage, measured by quasi-fixed costs over sales. (A similar plot results from using the book-to-market ratio.) The plot affirms the positive association between inflexibility and the sensitivity of returns to operating leverage or B/M. The positive association holds whether or not one controls for differences in asset risk, as seen by the sets of distinct symbols.

To be clear again, the theory is not inconsistent with there being a value effect driven by within-firm variation in productivity. That would be equivalent to the assertion that firms overall are characterized by the parameter sets that yield positive slopes (the top half of the graph), which correspond to ones with little reversibility of investment. However, if there is variation in flexibility, we would expect that conditioning on it should lead to variation in the value effect.

Figure 6: *Flexibility and the Effect of Operating Leverage*

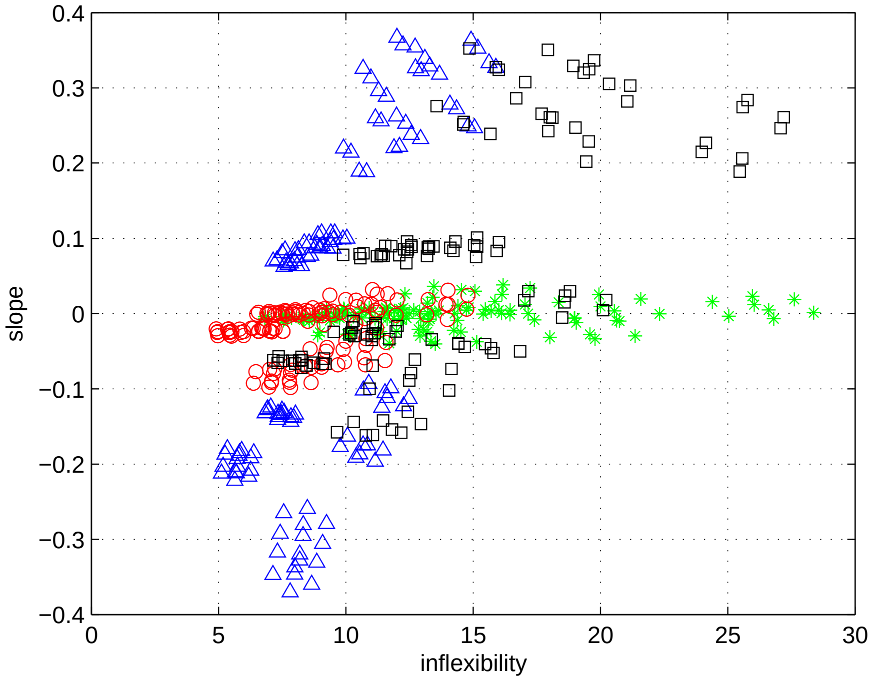

For each of the models described in Table 1, the slope of the graph of expected returns versus quasi-fixed costs over net sales is plotted here against a summary measure of that model's firm flexibility (the scaled range of its no-adjustment region). Cases with $\sigma = 0.55, \rho = -0.9$ are plotted as triangles. Cases with $\sigma = 0.25, \rho = -0.9$ are plotted as circles. Cases with $\sigma = 0.55, \rho = -0.1$ are plotted as squares. Cases with $\sigma = 0.25, \rho = -0.1$ are plotted as asterisks.

Summarizing, we have shown how different degrees of investment flexibility affect firms' risk/reward properties. The general lesson is that real-options contribute a downward-sloping function of operating leverage and assets-in-place contribute an upward-sloping one. Lower adjustment costs (higher flexibility) raises the influence of the former, but this does not lower total risk. Our theoretical analysis leads us to hypothesize that, if we can measure investment flexibility, it could show up empirically as a determinant of the response of expected returns to changes in profitability or operating leverage.

3 Measuring Inflexibility and Quasi-Fixed Costs

To take our model to the data, we need a way to differentiate firms according to their investment flexibility. We conjecture that the primary determinants of a firm's ability to adjust its scale derive from industry-wide features of physical and technological capital. Economic intuition suggests that industries differ as to what production inputs are acquired under long-term contracts, such as, some part of labor input, raw materials, and organization capital, and as to how easily productive capital is transformable. Hence we regard adjustment costs as a 'fact of life' for firms within an industry and propose time-invariant measures of inflexibility at the industry level. Within an industry, we can then assess each firm's profitability based on its expected, period-specific quasi-fixed production costs. Hence we attempt to measure time-varying quasi-fixed costs at the firm level. In this section, we present and discuss various measures of inflexibility and quasi-fixed costs, which we use in the tests of the next section.

To gauge an industry's inflexibility, we employ alternative classification schemes and data sources. Our primary proxies of industry inflexibility (*INFLEX*) are either directly derived from or closely linked to the model. That is, our estimates of *INFLEX* are based on cost stickiness and variability given that, according to the model, the observed range of profitability or quasi-fixed costs over sales increases with inflexibility. To this end, we build a measure of the median firm level range within an industry and a measure of the aggregate range of an industry. In addition, we supplement these measures of *INFLEX* by examining a resalability index for used industry capital and coefficient estimates of cost persistence obtained from industry panel regressions.

Our baseline inflexibility index (*INFLEX1*) is the standardized *median firm range* of operating costs (i.e., the sum of COMPUSTAT's costs of good sold, *COGSQ*, and, if available, selling, general, and administrative expenses, *XSGAQ*) over sales (i.e., *SALEQ*). More specifically, for each firm in an industry, the historical range of operating costs over sales is divided by the residual standard deviation from a regression of operating costs over sales on four of its own lags and a constant. The median firm range (i.e., *INFLEX1*) corresponds to the median value of these ranges across all firms in each of the 48 Fama and French (1997) industries.[12] Intuitively, an *INFLEX1* value of six can be roughly interpreted as the lower and upper boundaries (i.e., L and U) in the real options model being six standard deviations apart.

As seen in Table 2, *INFLEX1* ranges from 5.30 to 10.36. Thus, there is reasonable heterogeneity across industries, as also reflected by the standard deviation of about 0.88 relative to a median *INFLEX1* value of 6.95. While not all the rankings produced by this procedure have obvious causes in terms of industry features, the least flexible firms do include capital intensive manufacturing businesses, while several of the most flexible firms are notable users of outsourced capacity.[13]

---

[12] Following standard practice in the empirical asset pricing literature, we exclude banks (FF=44), insurance companies (FF=45), trading firms (FF=47) and utilities (FF=31) throughout.

[13] Note also from the third and fourth columns that none of the unexpected entries (e.g., coal) is large enough to have undue influence in the tests.

Table 2: *Industries with High and Low Inflexibility*

FF Code	Industry Description	Inflexibility	Number of Obs.	% Mkt. Cap.
	Panel A. Six industries with lowest inflexibility			
4	Beer & Liquor	5.30	13.84	2.61
29	Coal	5.52	6.38	0.15
16	Textiles	5.96	32.81	0.28
10	Apparel	6.03	57.86	0.51
15	Rubber & Plastics	6.05	41.86	0.27
35	Computer Mfg	6.22	157.30	3.95
	Panel B. Six industries with highest inflexibility			
32	Telecom/TV Networks	8.05	109.62	10.84
12	Medical Equipment	8.08	125.42	1.66
18	Construction	8.09	48.92	0.43
48	Unclassified	8.26	30.82	0.74
5	Tobacco Products	8.99	4.20	1.63
13	Pharmaceuticals	10.36	163.58	11.10

This table reports the six industries with the largest and smallest values of the *median firm range*, *INFLEX1*, which is an industry's median value of the firm level range of operating costs (COMPUSTAT's *COGSQ* and *XSGAQ*) over sales (i.e., *SALEQ*) and standardized by the residual volatility. The third and fourth columns show, for each industry, the average number of firm observations (Number of Obs.) and the average fraction of total market capitalization (% Mkt. Cap.) in each monthly cross-section of our sample period.

As a second proxy of inflexibility based on the model's notion of scaled range, we construct *INFLEX2* as the standardized *industry range*. That is, we compute industry aggregate cost, sales, and assets by summing over all quarterly firm observations in COMPUSTAT, with each calendar quarter using any available firm quarter reported during that quarter. Industry operating costs and industry sales are standardized by industry assets (i.e., the industry's aggregate value of COMPUSTAT's *ATQ*). The industry range is then determined by the historical range of aggregate, standardized operating costs over sales divided by the residual standard deviation from a regression of operating costs on contemporaneous sales and four lags of operating cost and sales and a constant.[14]

The third index of inflexibility builds directly on inter-industry variation in the reversibility, namely industry level capital resalability. Balasubramanian and Sivadasan (2009) define a capital resalability index ($RESAL$) as the fraction of total capital expenditure in an industry accounted for by purchases of used (as opposed to new) capital, computed at 4-digit SIC level. These authors construct their index using detailed data on both new and used capital expenditures collected and published by the U.S. Census Bureau in 1992. Intuitively, ability to re-sell physical assets is akin to a disinvestment option at the firm level. Indeed, in the model, flexibility is most closely tied to P_U, the real liquidity of physical capital. In industries where capital is firm-specific, there will be a less active secondary market in used capital, and the index will be low. Moreover Balasubramanian and Sivadasan (2009) show that the index is a significant factor in explaining other industry traits associated with greater reversibility. Based on this insight from the industrial organization literature, we define our third measure of inflexibility as *INFLEX3* = 1 - *RESAL*. This measure is time-invariant and is only available for a restricted sample of manufacturing firms (i.e., SIC codes 2000–3999).

As a fourth measure of cross-industry cost stickiness, we use ordinary-least-squares coefficient estimates from industry-by-industry panel regressions of firm-level operating costs scaled by assets on a moving average of four of its own lags and contemporaneous sales scaled by assets.[15] The alternative inflexibility proxies are mainly used to validate the findings for our baseline measure, *INFLEX1*, in various robustness checks in the next section and hence we omit further details of their construction.

Finally, we need to measure the firm-specific state variable: capital scaled by productivity. While increases in this variable generate increasing operating leverage, we do not attempt to measure operating leverage (which goes to infinity for unprofitable firms). Instead, we construct empirical counterparts of the ratio of quasi-fixed costs over sales,

---

[14] The correlation between *INFLEX1* and *INFLEX2* equals 0.2714, with a p-value of 0.0620.

[15] Instead of scaling by assets, we have also verified that weighted-least-squares estimates of this regression model provide similarly results when weighing by the reciprocal value of assets.

QFC/S, which increases monotonically increasing with the model's state variable.

Using quarterly COMPUSTAT data for the 1975–2009 period, we obtain annual, firm-level estimates of *QFC/S* by running five-year, rolling-window regressions of operating costs on their first lag and contemporaneous sales. The measure of *QFC/S* in the year following the 5-year estimation period equals the sum of regression intercept and predicted operating costs, scaled by sales. For inclusion in the sample, we require that quarterly growth rates in assets, costs, or sales lie inside the $[-75\%, +75\%]$ interval and that rolling-window regressions are based at least 10 observations. In robustness tests in the next section, we analyze the importance of the two components of *QFC/S* by using only the regression intercept scaled by sales as a measure of quasi-fixed costs. In another alternative specification, we reduce the noisiness of *QFC/S* estimates by increasing the number of observations from 10 to 15 for every 5-year window.

4 Empirical Evidence

Based on measures of inflexibility and quasi-fixed costs, we now examine whether the average sensitivity of returns to operating leverage depends on scale adjustment frictions. That is, the model's testable implication is that the strength of the relation between unprofitability (or quasi-fixed costs over sales) and stock returns increases with inflexibility. That is, if inflexibility is high (low), then expected stock returns increase (decrease) with QFC/S. This is a prediction about an interaction effect. The model does not have predictions about the marginal effect of either variable alone.

To start, to gauge the economic magnitude of the hypothesized effect, we consider the returns to portfolios formed based upon sorts on the two variables. Each month, we assign stocks into five quintiles based on two measures of industry inflexibility: median firm range (*INFLEX1*) and industry range (*INFLEX2*). We intersect these quintiles with a second independent sort of firms into quintiles according to their estimated quasi-fixed costs over sales. After assignment to portfolios, stocks are held for one month. We calculate the monthly portfolio return as the equal-weighted average of the returns of all the stocks in a portfolio. Figure 7 presents the average monthly return from 1980 to 2009. The figure plots the variation in the quasi-fixed cost sort on the horizontal axis, with separate lines for the portfolios of each industry quintile.

Figure 7: *Portfolio Returns for Double-Sorts*

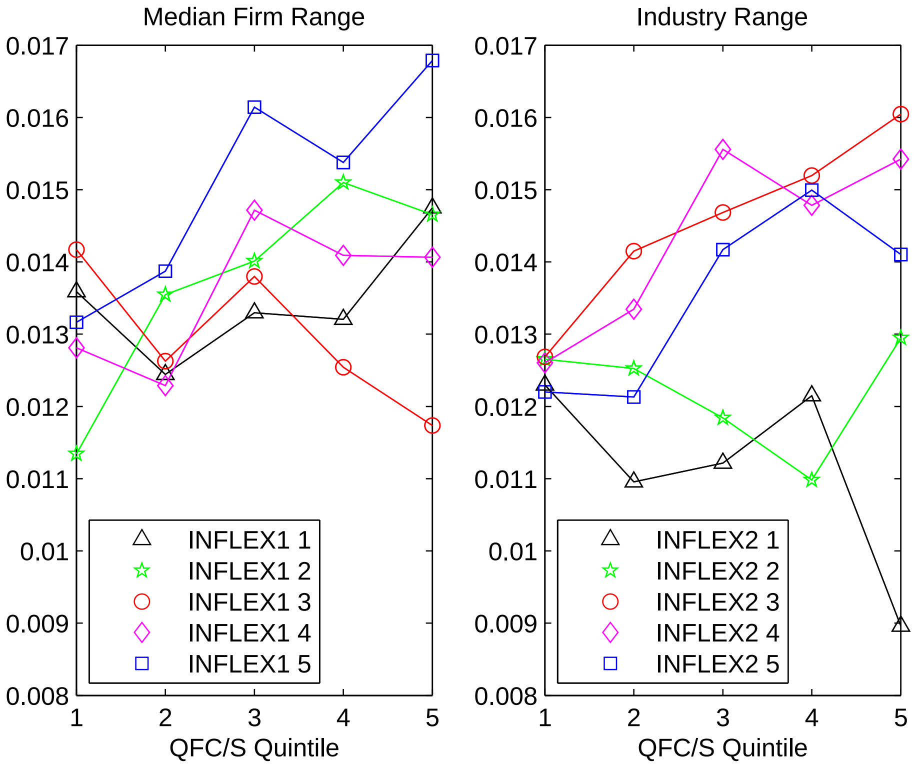

The figure shows the monthly profits from 25 portfolio strategies formed by independent sorts on firm-level quasi-fixed costs over sales and two measures of industry level inflexibility. The left panel measures inflexibility as the industry median of firm-level range of (*INFLEX1*); the right panel employs the industry range of costs over sales (*INFLEX2*). The sample period is January 1980 to December 2009.

The right-hand panel reveals a very strong interaction affect. With this measure of flexibility, there is a dramatic difference in slope between the more flexible industries (1 and 2) and the less flexible (4 and 5). Indeed, the resemblance to Figure 2 is remarkable. Economically, the difference in expected returns between most and least flexible industries for firms in the lowest quasi-fixed costs quintile is negligible, whereas in the highest quintile it is about 50 basis points per month, a very large number.

The left panel is less supportive. It remains the case that average returns increase most steeply with operating leverage in the least flexible industries (5), and that there is little evidence of a positive slope in the most flexible (1). The picture is muddled for the middle quintiles however. To some extent, this reflects limitations of the sorting methodology. The results below will show that this flexibility metric becomes highly informative once other determinants of expected return are controlled for.

To perform more formal tests, we follow the cross-sectional literature by testing the model's conditional return implications using standard Fama and MacBeth (1973) regressions.[16] In this context, the model says that the slope coefficient of an interaction term between inflexibility and quasi-fixed costs over sales should be positive and significant.

We carry out the tests using the intersection of the monthly stock returns from CRSP and quarterly COMPUSTAT accounting data for every month from January 1980 to December 2009. The baseline results are shown in Table 3. The first regression in Panel A displays the individual effect of the median firm range (*INFLEX1*) and quasi-fixed costs over sales (*QFC/S*) on expected stock returns, where *QFC/S* is winsorized at the 1% level. Neither variable is significant.

The second specification includes the interaction term (*INTER*). As predicted by the theory, the coefficient estimate is positive, but not significant. These specifications, however, do not control for other cross-sectional determinants of expected returns. The model's predictions apply to the incremental effect of inflexibility and quasi-fixed costs with other characteristics held constant. That is, heterogeneity in priced fundamental risk ($\rho\,\sigma$) and heterogeneity in financial leverage may affect the cross-sectional relationship.

In specification (4) of Panel A, we include standard control variables, namely, reversal (*R01*), momentum (*R12*), book-to-market ratio (*BM*), market leverage (*ML*), and size (*SZ*).[17] With this specification, the interaction term, *INTER*, increases greatly in

---

[16] As discussed in Section 2, the model's exact expression for expected returns is not expressible in closed form as a function of the parameters. This makes direct structural estimation infeasible, and may also rely too much on what remains a fairly stylized model.

[17] The variable *R01* is the stock return over the previous month; *R12* is the stock return over the 11 months preceding the previous month; *BM* denotes the log of the ratio of book value of equity to market value of equity; *ML* is the log of the market leverage ratio defined as book value of long-term

magnitude and statistical significance.

Furthermore, examining the coefficients in terms of economic significance in Panel A reveals that the effect is strong. For a firm in an inflexible industry (one standard deviation above the mean of *INFLEX1*), when quasi-fixed costs over sales go from one standard deviation below its mean to one standard deviation above, expected returns decrease by 5 basis point per month. For a firm in a flexible industry (one standard deviation below *INFLEX1*'s mean), the decrease is 51 basis points per month. This is consistent with the model's implication that flexible firms exercise disinvestment options in bad states and thereby reduce exposure to priced risk.

In Panel B of Table 3 all variables are transformed into percentile ranks to diminish the possible influence of outliers. As in Panel A, the individual influence of *QFC/S* and now also of *INFLEX1* on returns is unreliable in the first row of Panel B, while their interaction term, *INTER*, in the second row obtains a remarkably high level of statistical significance for a purely accounting-based variable. Again, including the other controls (specification (4)) further increases the statistical significance.

Consider now the marginal effects in Table 3. Other than regression (4) in Panel A, there is not a consistently positive coefficient on inflexibility. This is not inconsistent with the model. (Recall Figure 5.) But it runs counter to the conventional wisdom that flexible firms are necessarily safer. Flexibility increases the value of investment and disinvestment options, which in turn have opposing effects on risk exposure.

By contrast, the marginal effect of quasi-fixed costs on returns appears significantly negative. In terms of the model, this suggests that the average firm in the economy is relatively flexible, and, in particular, has some ability to reverse investment. This finding casts doubt on the conjecture that irreversibility is the driving force behind the "value premium." Moreover, in comparing specification (3) in each panel with the corresponding specification (4), we observe that the coefficient estimates on *BM* are undiminished by the presence of our variables. Neither the unconditional inflexibility effect nor the conditional

---

debt divided by the sum of market value of equity and book value of long-term debt; and *SZ* is the log of the market value of equity.

Table 3: *Return Regressions for Median Firm Range (INFLEX1)*

Panel A. Winsorizing QFC/S at the 1st and 99th Percentile

	INFLEX1	*QFC/S*	*INTER*	*R01*	*R12*	*BM*	*ML*	*SZ*
(1)	0.0936	–0.0379						
	(1.87)	(0.23)						
(2)	0.8242	–0.8313	0.1139					
	(1.78)	(1.00)	(1.01)					
(3)				–4.5160	0.5384	0.4335	–0.0565	–0.1131
				(10.51)	(2.87)	(7.49)	(2.08)	(2.40)
(4)	0.1900	–1.9790	0.2426	–4.6090	0.5169	0.4479	–0.0575	–0.1251
	(4.23)	(2.30)	(2.04)	(11.05)	(2.82)	(8.27)	(2.29)	(2.78)

Panel B. Transforming Variables into Percentile Ranks

	INFLEX1	*QFC/S*	*INTER*	*R01*	*R12*	*BM*	*ML*	*SZ*
(1)	0.0007	0.0021						
	(0.53)	(0.73)						
(2)	–0.0010	–0.0322	0.0344					
	(0.63)	(2.37)	(2.56)					
(3)				–0.0205	0.0124	0.0170	–0.0062	–0.0072
				(7.98)	(3.85)	(7.74)	(4.08)	(2.02)
(4)	–0.0001	–0.0560	0.0566	–0.0210	0.0119	0.0174	–0.0061	–0.0070
	(0.06)	(4.64)	(4.79)	(8.47)	(3.88)	(8.54)	(4.42)	(2.07)

The table shows estimation results from monthly Fama-MacBeth regressions of returns on measures of inflexibility (*INFLEX1*), quasi-fixed costs over sales (*QFC/S*), and their product (*INTER*), as well as on controls for expected returns. The variable *R01* is the stock return over the previous month; *R12* is the stock return over the 11 months preceding the previous month; *BM* denotes the log of the ratio of book value of equity to market value of equity; *ML* is the log of the market leverage ratio defined as book value of long-term debt divided by the sum of market value of equity and book value of long-term debt; and *SZ* is the log of the market value of equity. In Panel A, *QFC/S* is winsorized at the 1% level. In Panel B, all variables are transformed into percentile rank form. The data are monthly observations from January 1980 through December 2009. The coefficients are multiplied by 100 and t-statistics are in parentheses.

(interaction) effect with quasi-fixed costs over sales significantly lowers the explanatory power of the book-to-market ratio, suggesting that the value effect is more likely driven by cross-firm differences in risk than by within firm variation caused by quasi-fixed costs.

Table 4 reports estimation results for alternative measures of inflexibility and quasi-fixed costs over sales. Specifications (1)–(4) respectively use median firm range (*INFLEX1*), industry range (*INFLEX2*), capital illiquidity (*INFLEX3*), panel-regression estimated industry cost persistence (*INFLEX4*). In Panel A, *QFC/S* relies on the baseline definition (i.e., sum of regression intercept and predicted costs from rolling window estimations divided by sales). Notably, all coefficient estimates for the interaction term are reliably positive. It worth noting that the average sample size is almost reduced by 2/3 in regression model (3) compared to the other ones since the capital illiquidity index is only available for manufacturing firms (SIC codes 2000–3999). Therefore, it is even more remarkable that *INFLEX3* interacts significantly with *QFC/S*.

Panel B of Table 4 studies the importance of the two parts of *QFC/S* by dropping the predicted cost component, which might arguably be closer to variable than to fixed costs. That is, *QFC/S* in this set of tests is the regression intercept from a 5-year rolling estimation window ending in the year prior to the return observation, divided by sales. For all proxies of inflexibility in Panel B, the estimated interaction effect of inflexibility and quasi-fixed costs is smaller and statistically weaker than in Panel A, which underscores the importance of using both components of quasi-fixed costs.

Our baseline results are also robust to several alternative measurements of our key proxies. In unreported estimations, we find similarly interaction results when we (i) scale quasi-fixed costs by assets instead of sales, (ii) estimate *QFC/S* in terms of ratios of costs over sales instead of levels of costs on sales, or (iii) reduce the noisiness of *QFC/S* estimates by increasing the number of observations from 10 to 15 for every 5-year window.

Taken together, the empirical evidence in this section is strongly supportive of the model's prediction that flexible firms exercise disinvestment options if their profitability declines (i.e., operating leverage increases) lowering their exposure to fundamental (priced) risk and reducing expected stock returns, whereas inflexible firms with fewer

Table 4: *Return Regressions for Alternative Inflexibility Measures*

Panel A. Baseline Definition of QFC/S

	INFLEX	*QFC/S*	*INTER*	*R01*	*R12*	*BM*	*ML*	*SZ*
(1)	–0.0001	–0.0560	0.0566	–0.0210	0.0119	0.0174	–0.0061	–0.0070
	(0.06)	(4.64)	(4.79)	(8.47)	(3.88)	(8.54)	(4.42)	(2.07)
(2)	–0.0006	–0.0114	0.0121	–0.0211	0.0118	0.0173	–0.0059	–0.0068
	(0.42)	(2.34)	(2.48)	(8.45)	(3.82)	(8.41)	(4.26)	(2.02)
(3)	–0.0016	–0.0144	0.0152	–0.0207	0.0121	0.0171	–0.0058	–0.0070
	(1.09)	(3.10)	(3.10)	(8.26)	(3.94)	(8.37)	(4.03)	(2.06)
(4)	0.0033	–0.0044	0.0046	–0.0210	0.0120	0.0178	–0.0059	–0.0069
	(2.52)	(1.82)	(1.89)	(8.46)	(3.91)	(8.92)	(4.33)	(2.04)

Panel B. Alternative Definition of QFC/S

	INFLEX	*QFC/S*	*INTER*	*R01*	*R12*	*BM*	*ML*	*SZ*
(1)	0.0018	–0.0478	0.0493	–0.0208	0.0120	0.0173	–0.0061	–0.0069
	(1.42)	(3.14)	(3.23)	(8.19)	(3.76)	(8.05)	(4.07)	(1.97)
(2)	0.0008	–0.0058	0.0073	–0.0208	0.0119	0.0172	–0.0060	–0.0067
	(0.64)	(1.14)	(1.41)	(8.17)	(3.71)	(7.95)	(3.96)	(1.93)
(3)	0.0028	–0.0148	0.0156	–0.0189	0.0083	0.0095	–0.0017	–0.0041
	(2.42)	(0.43)	(0.46)	(8.64)	(2.68)	(5.00)	(1.01)	(1.18)
(4)	0.0041	–0.0017	0.0033	–0.0208	0.0122	0.0179	–0.0057	–0.0066
	(3.17)	(0.69)	(1.38)	(8.20)	(3.82)	(8.63)	(3.95)	(1.88)

The table shows estimation results from monthly Fama-MacBeth regressions of returns on measures of inflexibility (*INFLEX*), quasi-fixed costs over sales (*QFC/S*), and their product (*INTER*), as well as on controls for expected returns. In each panel, specification (1) uses the median range (*INFLEX1*); specification (2) uses the industry range (*INFLEX2*); specification (3) uses the transformed capital resalability index (*INFLEX3*); and specification (4) uses the panel-regression estimate of industry cost persistence (*INFLEX4*). The variables *R01*, *R12*, *BM*, *ML*, and *SZ* are defined in the caption of Table 3. In Panel A, *QFC/S* is the sum of regression intercept and predicted operation costs from a 5-year rolling estimation window ending in the year prior to the return observation, divided by sales. In Panel B, *QFC/S* is the regression intercept from a 5-year rolling estimation window ending in the year prior to the return observation, divided by sales. All variables are transformed into percentile rank form. The data are monthly observations from January 1980 through December 2009. The coefficients are multiplied by 100 and *t*-statistics are in parentheses.

(or more costly) disinvestment options experience the reverse. The results underline the theoretical richness of the neoclassical framework and even helps explain cross-sectional variations of stock returns in practice.

5 Conclusion

Much insight about the cross-section of stock returns has emerged from viewing firms as being essentially equal *ex ante* but differing in their installed investment base and in their current production opportunities. We augment this class of models by examining the further cross-sectional implications of heterogeneity in firm flexibility, or adjustment costs. Of course, this is just one dimension along which real firms vary. The literature has, however, shown that flexibility is crucial in determining how operating risks translate into shareholder risks. In particular, the extreme case of irreversible capital, combined with quasi-fixed operating costs, implies strong variation in the equity risk premium with profitability (or book-to-market or Tobin's Q). This suggests that differences in flexibility may entail interesting differences in this relationship.

Indeed, in the context of simple, one-state variable, partial equilibrium model, we show that irreversibility (or very low recovery values of disinvested capital) is necessary for a positive average association between book-to-market and expected returns. Moderate reversibility implies the opposite. Whether or not the model implies an overall "value premium" then depends on the degree of reversibility in actual firms. When we interact measures of industry inflexibility with profitability (quasi-fixed cost over sales) in return regressions, we do find a significant positive association, as predicted by the model. But we do not find that these variables at all diminish the importance of book-to-market, suggesting that the value effect is more likely driven by cross-firm differences in fundamental risk than by within-firm risk variation caused by operating leverage.

References

Abel, Andrew B., Avinash Dixit, Janice C. Eberly, and Robert S. Pindyck, 1996, Options, the value of capital, and investment, *The Quarterly Journal of Economics* 111, 753–77.

Abel, Andrew B, and Janice C Eberly, 1996, Optimal investment with costly reversibility, *Review of Economic Studies* 63, 581–93.

Aguerrevere, Felipe L., 2009, Real options, product market competition, and asset returns, *Journal of Finance* 64, 957–983.

Balasubramanian, Natarajan, and Jagadeesh Sivadasan, 2009, Capital resalability, productivity dispersion, and market structure, *Review of Economics and Statistics* 91, 547–557.

Carlson, Murray, Adlai Fisher, and Ronald Giammarino, 2004, Corporate investment and asset price dynamics: Implications for the cross-section of returns, *Journal of Finance* 59, 2577–2603.

Chirinko, Robert S., and Huntley Schaller, 2009, The irreversibility premium, *Journal of Monetary Economics* 56, 390–408.

Cooper, Ilan, 2006, Asset pricing implications of nonconvex adjustment costs and irreversibility of investment, *Journal of Finance* 61, 139–170.

Fama, Eugene F., and Kenneth R. French, 1997, Industry costs of equity, *Journal of Financial Economics* 43, 153–193.

Fama, Eugene F., and James D. MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* 21, 607–636.

Gala, Vito D., 2006, Investment and returns, London Business School working paper.

Garca-Feijó, Luis, and Randy D. Jorgensen, 2010, Can operating leverage be the cause of the value premium?, *Financial Management* 39, 1127–1154.

Garlappi, Lorenzo, Tao Shu, and Hong Yan, 2008, Default risk, shareholder advantage and stock returns, *Review of Financial Studies* 81, 2743–2778.

Garlappi, Lorenzo, and Hong Yan, 2011, Financial distress and the cross section of equity returns, *Journal of Finance* 66, 789–822.

Gomes, Joao, Leonid Kogan, and Lu Zhang, 2003, Equilibrium cross section of returns, *Journal of Political Economy* 111, 693–732.

Guthrie, Graeme, 2010, A note on operating leverage and expected rates of return, *Finance Research Letters* forthcoming.

Li, Erica X. N., Dimitry Livdan, and Lu Zhang, 2009, Anomalies, *Review of Financial Studies* 22, 4301–4334.

MacKay, Peter, 2003, Real flexibility and financial structure: An empirical analysis, *Review of Financial Studies* 16, 1131–1165.

Novy-Marx, Robert, 2011, Operating leverage, *Review of Finance* 15, 103–134.

Ozdagli, Ali K., 2011, Financial leverage, corporate investment, and stock returns, *Review of Financial Studies* forthcoming.

Sagi, Jacob, Matthew Spiegel, and Masahiro Watanabe, 2009, Dynamic Corporate Capital Stocks: Cross-sectional and Inter-temporal Stock Return Patterns, Vanderbilt University working paper.

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

Appendix: Model Solution

The text describes the form of the firm's impulse control policy. When it is in the no-adjustment region the firm value, J, satisfies the equilibrium condition

$$\mathrm{E}[dJ/J] + \Pi/J - r = -\mathrm{Cov}[dJ/J, d\Lambda/\Lambda].$$

which implies the partial differential equation

$$[J_\theta \; \theta \; \mu - J_K \; \delta \; K + \frac{1}{2} J_{\theta\theta} \theta^2 \; \sigma^2] + \Pi - rJ + [\rho \; J_\theta \; \theta \; \sigma \; \sigma_\Lambda] = 0. \tag{A.1}$$

To verify homogeneity, we guess the solution form $J = \theta \; V(K/\theta)$ and use

$$\begin{aligned} J_K &= V' \\ J_\theta &= V - (K/\theta) \; V' \\ J_{\theta\theta} &= -(K/\theta^2) \; V' + (K/\theta^2) \; V' + (K^2/\theta^3) \; V'' = (K^2/\theta^3) \; V''. \end{aligned}$$

The PDE then becomes an ODE in $V(Z)$ with $Z = K/\theta$:

$$\frac{1}{2} Z^2 \; \sigma^2 \; V'' - [\mu + \rho \; \sigma \; \sigma_\Lambda + \delta] \; Z \; V' + [\mu + \rho \; \sigma \; \sigma_\Lambda - r] \; V + [Z^\gamma - mZ] = 0. \tag{A.2}$$

Note that the risk neutral drift of the productivity process is $\mu^{RN} \equiv \mu + \rho \; \sigma \; \sigma_\Lambda$. A regularity condition of the problem is that $\mu^{RN} < r$.

In terms of the re-scaled variable Z and the function V, the task is to choose points G, L, U, H on the positive Z axis to maximize V. Absence of arbitrage imposes the two value matching conditions (VMCs):

$$V(G) = V(L) + F_L L^\gamma + P_L \, (G - L) \tag{A.3}$$

and

$$V(H) = V(U) + F_U U^\gamma + P_U \, (H - U) \, . \tag{A.4}$$

The first equation requires that the post-investment value of the firm is the pre-investment value plus the funds injected. The second imposes the same for pre- and post- disinvestment (note $H - U < 0$). Given these, functionally differentiating with respect to the barrier positions, yield the smooth-pasting conditions (SPCs) as necessary conditions of optimality. These are:

$$\begin{aligned} V'(L) &= -\gamma\ F_L L^{\gamma-1} + P_L, && \text{(A.5)} \\ V'(G) &= P_L, && \text{(A.6)} \\ V'(U) &= -\gamma\ F_U U^{\gamma-1} + P_U, && \text{(A.7)} \\ V'(H) &= P_U. && \text{(A.8)} \end{aligned}$$

The solution to (A.2) is well known: it is the sum of the general form of solution to the homogenous version (without the Π terms) and a particular solution having the same form as the Π terms. This yields

$$V(Z) = A\ Z^\gamma - S\ Z + D_N\ Z^{\lambda_N} + D_P\ Z^{\lambda_P} \tag{A.9}$$

where

$$A = \frac{1}{r + \gamma\delta + (\gamma-1)\mu^{RN} - \frac{1}{2}\gamma(\gamma-1)\sigma^2}$$

$$S = \frac{m}{(r+\delta)}$$

and

$$\lambda_{P,N} = \frac{b \pm \sqrt{b^2 + 2(r-\mu^{RN})\sigma^2}}{\sigma^2}$$

and $b = (\mu^{RN} + \delta + \frac{1}{2}\sigma^2)$. Here D_N and D_P are two additional free parameters.

When (A.9) is plugged into each of the SPCs and VMCs, the result is a system of six equations in G, L, U, H, D_N, and D_P. The system is linear in the last two, given the first four. But the nonlinearity in the first four renders numerical solution necessary. Solving the VMCs yields

$$\begin{aligned} D_N &= \frac{1}{\Delta}\Big[(H^{\lambda_P} - U^{\lambda_P})(A(G^\gamma - L^\gamma) - S(G-L) - F_L L^\gamma - P_L(G-L)) \\ &- (G^{\lambda_P} - L^{\lambda_P})(A(H^\gamma - U^\gamma) - S(H-U) - F_U U^\gamma - P_U(H-U))\Big] \end{aligned} \tag{A.10}$$

and

$$D_P = \frac{1}{\Delta}\Big[(G^{\lambda_N} - L^{\lambda_N})(A(H^\gamma - U^\gamma) - S(H-U) - F_U U^\gamma - P_U(H-U))$$

$$- \ (H^{\lambda_N} - U^{\lambda_N})(A(G^{\gamma} - L^{\gamma}) - S(G - L) - F_L L^{\gamma} - P_L(G - L))\Big] \quad \text{(A.11)}$$

where

$$\Delta = (G^{\lambda_P} - L^{\lambda_P})(H^{\lambda_N} - U^{\lambda_N}) - (G^{\lambda_N} - L^{\lambda_N})(H^{\lambda_P} - U^{\lambda_P}).$$

As discussed in the text, another facet of the problem is the abandonment option. If the solution found by the above procedure does not yield an everywhere positive firm value (which can happen, for example, if P_U is very negative), then it is not consistent with limited liability. In that case, the system is re-solved with the boundary conditions $V(U) = 0$ and $V'(U) = 0$ replacing (A.4) and (A.7).

Based on (A.9)–(A.11), the expected excess return can be written

$$\begin{aligned} EER(Z) &= \pi_\theta \, \frac{AZ^{\gamma}(1-\gamma) + D_N Z^{\lambda_N}(1-\lambda_N) + D_P Z^{\lambda_P}(1-\lambda_P)}{V(Z)} \\ &= \pi_\theta \, \frac{\widehat{V(Z)}}{V(Z)} \, \frac{AZ^{\gamma}(1-\gamma) + D_N Z^{\lambda_N}(1-\lambda_N) + D_P Z^{\lambda_P}(1-\lambda_P)}{\widehat{V(Z)}} \end{aligned}$$

where we define $\widehat{V(Z)} \equiv V(Z) + SZ = AZ^{\gamma} + D_N Z^{\lambda_N} + D_P Z^{\lambda_P}$, which is the value of the firm excluding the liability due to fixed costs. Thus we have[18]

$$EER(Z) = \pi_\theta \, \frac{\widehat{V(Z)}}{V(Z)} \, \left[(1-\gamma)w_A(Z) + (1-\lambda_N)w_G(Z) + (1-\lambda_P)w_C(Z)\right].$$

where $w_A + w_C + w_G = 1$. If the firm's real options were worthless, the only variation in the risk premium would come from the term $\frac{\widehat{V(Z)}}{V(Z)} = \left(1 + \frac{SZ}{V(Z)}\right)$, which represents the risk amplification of operating leverage. (By contrast, the weights w_A, w_C, w_G do not depend on the value of the fixed-cost liability.) The risk contribution of assets in place is given by $(1-\gamma)w_A(Z) \ (\widehat{V(Z)}/V(Z))$. Clearly w_A is increasing in Z. The operating leverage term is as well since the derivative of $1 + SZ/V$ has the same sign as that of Z/V, which is $(V - ZV')/V^2 > 0$ (c.f., property (A) in Section 2). It follow that the total risk premium contribution of assets in place is increasing in Z as claimed in the text.

By the same reasoning, the contraction option's contribution to the risk premium, $(1-\lambda_P)w_C(Z) \ (\widehat{V(Z)}/V(Z))$, is increasingly negative (since $\lambda_P > 1$ and $w_C(Z)$ is in-

---

[18] We thank Ali Ozdagli for suggesting this decomposition.

creasing). The remaining claim in the text is that the growth option term is *decreasing* in Z. This term is equal to $(1-\lambda_N)D_N Z^{\lambda_N}/V$ and $(1-\lambda_N) > 0$ and $D_N > 0$. The sign of the derivative is thus the sign of $\lambda_N V - ZV'$. Since $V > 0$ by limited liability, and $\lambda_N < 0$, a sufficient condition for a negative derivative is simply $V' > 0$. But this is implied by condition (B) in Section 2.