2019

What Can We Learn from Stock Prices? Cash Flow, Risk and Shareholder Welfare

Joshua Mitts
Columbia Law School, jmitts@law.columbia.edu

What Can we Learn from Stock Prices?: Cash Flow, Risk, and Shareholder Welfare

by

Joshua Mitts*

Price is expected cash flows discounted at the risk-free rate plus an additional discount for risk exposure. Price equivalency does not always imply welfare equivalency: shareholders are not necessarily indifferent between a price increase of $1 from higher cash flows and the same $1 increase from lower risk exposure. Even in complete markets, if managers enjoy private benefits of control, the social planner may prefer lower risk exposure to a price-equivalent increase in firm value from greater investor protection. This has implications for event studies, the trade-off between principal costs and agency costs, and the link between macroeconomic risk and corporate governance.

Keywords: securities law, firm value, asset prices, shareholder welfare

JEL classification code: D53, G12, G34, K22

1 Introduction

Prices convey information. Hayek (1945, p. 526) put it this way: prices "coördinate the separate actions of different people in the same way as subjective values help the individual to coördinate the parts of his plan." Stock prices, in particular, matter a great deal in corporate and securities law. Event studies, which measure statistically significant changes in stock prices, are widely used by investors and courts to infer the effect of an event on the value of a firm (Bhagat and Romano, 2002).

This article asks a basic question: what can we learn from stock prices? It is a tautology to say that price reflects value: after all, buyers will not pay more for an asset than what that asset is worth to them. But *value* does not imply *cash flow*: a buyer may happily pay $1 for an asset that never pays off $1. That is because assets that pay off when investors need money are more valuable to those investors than assets that pay off when they are wealthy. Stock prices reflect not only the expected future cash flows of a firm, but also the extent to which those cash flows

* Associate Professor of Law, Columbia Law School, New York, USA. I am grateful to my discussants, Stephen Choi and Erik Theissen, for very thoughtful comments. Many thanks to Ian Ayres, Lucian Bebchuk, Patrick Bolton, Ronald Gilson, Larry Glosten, Jeffrey Gordon, Zohar Goshen, Merritt Fox, Gur Huberman, Robert J. Jackson, Jr., Wei Jiang, Edward Morrison, and Eric Talley for their helpful feedback.

Journal of Institutional and Theoretical Economics 175, 178–195 – ISSN 0932-4569
DOI: 10.1628/jite-2019-0009 –

serve as a kind of implicit insurance for investors, paying off precisely when they need money. Put differently, investors especially value cash flows that *smooth consumption* by making hard times less painful.

I make a simple claim: the distinction between increasing expected cash flows and the insurance function that those cash flows serve matters for welfare. A $1 increase in the price of a firm can have different welfare implications depending on what is driving it. I outline two conditions under which this distinction matters. The first is *imperfect risk-sharing*. It is an axiom of finance theory that the price of an asset is its expected cash flows discounted at the risk-free rate plus an additional discount for risk exposure. This latter component is the covariance of the asset's cash flows with the marginal utility of consumption, so-called *systematic risk* (Cochrane, 2005). In an endowment economy with complete risk-sharing, shareholders are indifferent between a $1 increase in expected cash flows and a $1 reduction in risk exposure. This is a consequence of the first welfare theorem (LeRoy and Werner, 2014, p. 164).

But perfect risk-sharing is an unrealistic ideal. There are all kinds of risks – especially those correlated with macroeconomic volatility, like losing a job or customers in a downturn – that cannot be fully insured. When risk-sharing is imperfect, investors are not indifferent between a $1 increase in the value of a firm driven by an increase in expected cash flows and a $1 increase in the value of the firm driven by a reduction in exposure to systematic risk. In some cases, investors may prefer the latter to the former. One reason, explored further in this article, is that prices are derived from marginal, not total, utilities. An agent may be willing to pay the same price for a reduction in risk and an increase in expected cash flow, even though the former yields a higher total utility than the latter.

In short, *price equivalency* does not imply *welfare equivalency*. Shareholders sometimes prefer a reduction in risk that yields a $1 increase in the value of the firm to an increase in expected cash flows that yields the same $1 increase. The second half of this article shows that a similar result may obtain even when risk-sharing markets are complete. I examine the classical setting where the firm is controlled by a manager who enjoys noncontractible private benefits of control (Jensen and Meckling, 1976). The intuition is simple: controlling and outside shareholders both prefer risk reduction, but they have opposite preferences as to the private benefits of control. A reduction in risk that is price-equivalent to an increase in investor protection will thus always be preferred by the controlling shareholder, while outside shareholders are indifferent (after all, that is exactly what price equivalency implies in complete markets).[1]

Section 2 provides sufficient conditions for risk reduction to yield a greater welfare increase than a price-equivalent increase in cash flows. Section 3 presents a

[1] These results are drawn from Mitts and Mansouri (2018), which provides a detailed exposition of the trade-off between disagreement and private benefits. Due to space constraints, I am only able to present basic results under constant absolute risk aversion (CARA) utility; in that project, we consider incomplete markets and alternative preferences in greater detail.

dynamic, continuous-time model that extends these intuitions to production and investment decisions. Section 4 discusses three implications of these findings for the event study methodology, the trade-off between principal costs and agency costs, and the link between macroeconomic risk and corporate governance.

2 *Exchange Economy*

I briefly review the simple general-equilibrium setting in Weil (1992) to illustrate the core intuition. Consider a two-period economy where consumers have time-additive preferences over consumption:

$$U(c_1,c_2) = u(c_1) + u(c_2).$$

A continuum of identical consumers is endowed with w_1 units of a consumption good at time 1. Each consumer receives a random endowment of income at time 2 denoted by y, which has cumulative distribution function F_y for every agent. As in Weil (1992), $\int y\,dy = \bar{y}$, a known constant, so there is no aggregate time-2 risk. Moreover, the time-2 income stream is nondiversifiable. There are a variety of reasons why labor market risks cannot be pooled, including moral hazard and asymmetric information. Asset markets are thus ex post incomplete: the risk represented by y is not insurable.

There is a risky asset that pays a random dividend d for every agent. Dividend risk is thus systematic. Letting x denote the amount of wealth invested in the risky asset, the consumer's problem at time 1 is given by

$$\max_x u(c_1(x)) + E[u(c_2(x))]$$

subject to

$$c_1 + px \leq w_1, \quad c_2 \leq dx + y, \quad c_1, c_2 \geq 0.$$

Under the standard regularity conditions, the budget constraints hold with equality:

$$\max_x u(w_1 - px) + E[u(dx + y)],$$

which yields the standard first-order condition for the price of the risky asset:

$$-u'(w_1 - px)p + E[u'(dx + y)d] = 0,$$
$$p = E\left[\frac{u'(dx + y)}{u'(w_1 - px)} d\right].$$

In equilibrium, the market-clearing condition is that supply equals demand. As agents are identical, the equilibrium does not admit trading. Normalizing the supply of the risky security to 1, we have $\int x\,dx = 1$. With identical agents, $x = 1$. Letting

$$m = \frac{u'(d + y)}{u'(w_1 - p)},$$

we have the familiar form (Cochrane, 2005):

$$p = E[md],$$

and by the definition of covariance,

$$p = E[m]E[d] + \text{cov}(m, d).$$

It is straightforward to see that $E[m] = 1/R_f$, where R_f is the risk-free rate. Thus,

$$p = \frac{E[d]}{R_f} + \text{cov}(m, d),$$

that is, the price of an asset is its expected payoff discounted at the risk-free rate plus a risk discount, which is the covariance of the asset payoff with the marginal utility of equilibrium consumption at time 2.

Cash Flow versus Risk Reduction. Suppose that the social planner can intervene either by increasing the expected cash flow $E[d]$ or by increasing its covariance with the marginal utility of time-2 consumption, specifically, $\text{cov}(m,d)$. Might the social planner prefer to reduce $\text{cov}(d, y)$ rather than $E[d]$, or vice versa? It is most interesting to make this comparison by holding fixed the price of the asset.

To see how the same price might have different welfare implications, consider the simple case of quadratic utility, which admits parsimonious analytical forms:

$$U(c) = Kc - \frac{1}{2}c.$$

Marginal utility is simple:

$$U'(c) = K - c.$$

Normalizing time-1 consumption to $U'(w_1 - p) = 1$, recall that the agent's time-2 consumption is $d + y$. Thus $m = K - (d + y)$, and the price of the asset is

$$p(d) = (K - E[d + y])E[d] + \text{cov}(K - (d + y), d),$$

or

$$p(d) = (K - E[d + y])E[d] - \text{var}(d) - \text{cov}(d, y).$$

Consider the simple case where $y \sim N(1,1)$ and $d = \beta(y - \bar{y}) + \varepsilon$ with $\varepsilon \sim N(1,1)$ and $\text{cov}(\varepsilon, y) = 0$. Thus $E[y] = E[d] = 1$, $\text{var}(d) = 1 + \beta^2$, and $\text{cov}(d, y) = \beta$.

Begin by considering the price of the asset when it receives an exogenous increase in cash flows. Denote a new asset with payoff $\tilde{d} = d + a$, where $a > 0$ is a constant. Of course, both $\tilde{d}$ and d have the same expected value, because the constant payoff a is simply added in every state: $E[\tilde{d}] = E[d] + a$. The price of $\tilde{d}$ is

$$p(\tilde{d}) = (K - E[d + a + y])E[d + a] - \text{var}(d) - \text{cov}(d, y).$$

Substitution yields

$$p(\tilde{d}) = (K-2-a)(1+a) - 1 - \beta^2 - \beta.$$

Now, consider another modification to d, but this time instead of adding a, simply change the covariance of its payoff with labor income. Let b denote this modified covariance and $\hat{d}$ the payoff: $\hat{d} = b(y-\bar{y}) + \varepsilon$. We obtain by simple substitution

$$p(\hat{d}) = K - 3 - b^2 - b.$$

Price Equivalency and Welfare. Price equivalency means $p(\tilde{d}) = p(\hat{d})$, that is,

$$(K-2-a)(1+a) - 1 - \beta^2 - \beta = K - 3 - b^2 - b.$$

A simple application of the quadratic formula yields the following solution for b:[2]

$$b^* = \frac{1}{2}\left(-1 \pm \sqrt{4\beta^2 + 4\beta + 1 + 12a + 4a^2 - 4aK}\right).$$

Suppose, for example, the "original" covariance $\beta = 0$ and $K = 100$. Consider $a = 0.001$. We obtain two possible solutions: $b^* = -0.891153$ and $b = -0.108847$. Both yield a positive price $p(\hat{d}) = 97.097$, which is equal to $p(\tilde{d})$ by construction. That is, those values of b^* are price-equivalent to $a = 0.001$.

The statement that risk reduction yields greater welfare than a cash-flow increase implies $E[U_b(\hat{d}+y)] - E[U_a(\tilde{d}+y)] > 0$. Recall that the expected utility is given by

$$E[U(d+y)] = K(E[d] + E[y]) - \frac{1}{2}\left[\mathrm{var}[d] + \mathrm{var}[y] + 2\mathrm{cov}(d,y) + (E[d] + E[y])^2\right].$$

Again, suppose for simplicity that $K = 100$ and $\beta = 0$. Substituting yields

$$E[U_b(\hat{d}+y)] = 200 - \frac{1}{2}[2 + 2b + 4],$$
$$E[U_a(\tilde{d}+y)] = 200 + 100a - \frac{1}{2}\left[2 + (2+a)^2\right].$$

It follows that risk reduction yields a higher expected utility than a price-equivalent increase in cash flows iff

$$\frac{1}{2}\left(a^2 - 196a \pm \sqrt{1 + 4a^2 - 97a} + 1\right) > 0.$$

[2] This is a sufficiency condition. A real-valued b^* may not exist. In addition, there can be multiple covariances that yield the same price as a change in cash flow: price equivalency is nonunique.

Figure
Cash-Flow Increase versus Price-Equivalent Risk Reduction

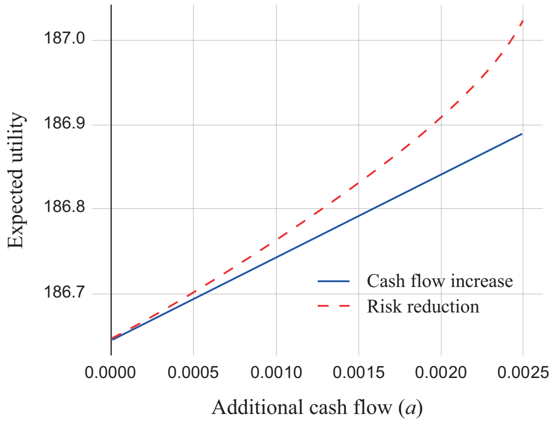

One interval satisfying this inequality is $a \in (0,0.0025]$. For any a in that interval, a price-equivalent change in covariance (that does not increase the expected payoff of the asset) yields a greater welfare increase than increasing the expected payoff of the asset by a (and not changing its covariance). The figure simulates $E[U_b(\hat{d}+y)]$ and $E[U_a(\tilde{d}+y)]$ over this interval to illustrate the welfare difference between risk reduction and cash flow.

As the figure shows, a reduction in risk yields a higher expected utility than a price-equivalent increase in cash flows for $a \in (0,0.0025]$. Of course, these magnitudes are small in this illustrative example – they simply show sufficient conditions for a price-equivalent reduction in risk to yield greater welfare than an increase in cash flow. The next section presents sufficient conditions for a similar result to obtain under a dynamic model with perfect diversification of idiosyncratic risk.

3 Complete Markets and Investment

The preceding section considered the difference in shareholder welfare between risk reduction and increasing cash flow in a simple, static endowment economy where the payoff of the firm is an exogenous random variable. I now consider a dynamic model where the value of the firm is derived from endogenous production and investment decisions. In this setting, the inability to contract over a controlling shareholder's diversion of cash flows for private benefits (see Jensen and Meckling, 1976) leads to a similar wedge between the welfare implications of risk reduction and cash-flow increase, even when these yield an equivalent price of the firm.

Crucially, a divergence between price equivalency and welfare equivalency arises *even when markets are complete*, i.e., there is perfect risk-sharing. The model is detailed in Mitts and Mansouri (2018); here, I sketch a few intuitions from that project under a simple setting where agents have constant absolute risk aversion (CARA) utility. There, we consider incomplete markets as well as constant relative risk aversion (CRRA) utility, which is mathematically more involved but yields similar intuitions.

Production. We follow standard models in this literature, supposing that the capital stock of the firm consists of tangible and intangible assets, which evolve deterministically in continuous time:

$$dK_t = (I_t - \delta_K K_t)dt,$$

where K_t denotes the capital stock of the firm, I_t denotes real investment, and $\delta_K \in (0,1)$ denotes the capital depreciation rate. In this continuous-time framework, this implies that over a small interval $(t, t+dt)$, capital stock increases by the investment made at time t (given by I_t), and the existing capital depreciates at a rate δ_K.

As in Lan, Wang, and Yang (2012), the firm is able to generate income by utilizing its assets for production, i.e., following the standard formulation $A_t K_t$. The firm's productivity A_t is subject to shocks, which constitute the only source of risk for the firm in this economy. We model productivity shocks as Brownian motion with mean $\mu_A > 0$ and volatility $\sigma_A > 0$ and suppose that the consumption of private benefits decreases expected productivity. This reflects the intuition that private benefits redirect scarce capital from productive uses. In short, productivity shocks to capital evolve according to the stochastic process

$$dA_t = (\mu_A - s_t)dt + \sigma_A\left(\rho_A dZ_t^M + \sqrt{1-\rho_A^2}dZ_t^A\right),$$

where μ_A is the expected productivity, s_t is the private benefits consumed at time t, and dZ_t^M and dZ_t^A are independent standard Brownian motions. As explained below, dZ_t^M is the source of tradable risk in the economy, i.e., risk that may be fully hedged by acquiring a risky security. dZ_t^A is idiosyncratic, firm-specific productivity risk that is orthogonal by construction to market risk. The correlation coefficient ρ_A reflects the extent to which the firm's productivity shock is correlated with tradable market risk.

Investment. Real investment imposes adjustment costs that are homogeneous of degree one in both investment and capital (see Hayashi, 1982). Letting $\Phi(I,K)$ denote the adjustment cost function, we have

$$\Phi(I,K) = \phi(i)K = \frac{\theta_i}{2}i^2K.$$

We model legal constraints on the consumption of private benefits as a convex investor protection cost that the controlling shareholder must pay for each private benefit extracted per unit of capital:

$$\psi(s) = \frac{\theta_s}{2} s^2,$$

where s denotes the private benefit per unit of capital and θ_s denotes the strength of a legal regime protecting investors from diversion of cash flows by the controlling shareholder (see La Porta et al., 2002, Johnson et al., 2000, and Stulz, 2005).

Letting Y_t denote the cumulative dividend to outside shareholders, we observe that Y_t follows a stochastic process given by

$$dY_t = (K_t \mu_A - I_t - \Phi(I_t, K_t) - s_t K_t)dt + \sigma_A \left(\rho_A dZ_t^M + \sqrt{1-\rho_A^2} dZ_t^A\right) K_t.$$

To see how this is derived, recall that the firm's income from the productive use of assets is given by $K_t A_t$. The above expression arises from considering the change in capital and productivity over the infinitesimal time interval $(t, t+dt)$, during which outside shareholders receive the gross income generated by the productive use of capital (the current capital K_t multiplied by the productivity shock dA_t), and adding the change in capital dK_t over the same window.

Similarly, letting M_t denote the cumulative ownership payoff to the controlling shareholder, we observe that M_t follows a stochastic process given by

$$dM_t = \alpha dY_t + s_t K_t dt - \psi(s) K_t dt,$$

where α denotes the size of the controlling shareholder's block (exogenously given), $s_t K_t$ is the per-period private benefit, and $\psi(s) K_t$ is the investor protection cost.

We allow for risk hedging via the financial markets. There is a risk-free asset yielding a return of r, and a risky asset whose return follows the stochastic process

$$dR_t^M = \mu_M dt + \sigma_M dZ_t^M,$$

where dZ_t^M is standard Brownian motion. The market price of risk is defined as

$$\eta = \frac{\mu_M - r}{\sigma_M}.$$

Complete Markets. In Mitts and Mansouri (2018), we consider the more general settings of incomplete markets. Here, for simplicity, we complete markets by introducing a traded financial asset that is perfectly correlated with the idiosyncratic productivity shock. These risks can be fully diversified away at no premium, which implies that the idiosyncratic hedge asset follows the stochastic process

$$dR_t^I = r dt + \sigma_I dZ_t^A,$$

where σ_s is the volatility parameter and Z_t^A is the same Brownian motion driving the nonmarket risk in the firm's productivity process.

The total wealth of the controlling shareholder is the sum of the ownership payoff and the return on the risky asset and hedge contract. Let X_t denote this total wealth, and Π_t denote the proportion invested in the risky asset, which implies that $X_t - \Pi_t$ is invested in the risk-free asset. Let Ω_t denote the amount of agent's wealth allocated to the idiosyncratic hedge. Then the controlling shareholder's total wealth follows the stochastic process

$$dX_t = r(X_t - \Pi_t)dt + (\mu_M dt + \sigma_M dZ_t^m)\Pi_t + \Omega_t \sigma_s dZ_t^A - C_t dt + dM_t.$$

The controlling shareholder chooses consumption, investment, portfolio allocation, and private benefits to maximize the expected discount value of lifetime consumption flows, which yields the standard formulation:

$$\max_{C_t,\Pi_t,I_t,s} E\left[\int_0^\infty e^{-\xi t} U(C_t)dt\right],$$

where $\xi > 0$ denotes the discount rate. All agents have constant absolute risk aversion (CARA) utility with

$$U(C_t) = -\frac{e^{-\gamma C_t}}{\gamma}.$$

Model Solution. Letting $F(K,X)$ denote the controlling shareholder's value function, by Ito's lemma the Hamilton–Jacobi–Bellman equation (HJB) is given by

$$\begin{aligned}\xi F(X,K) = \max_{C,\Pi,\Omega,I,s} U(C) + (I - \delta_K K)F_K + \Bigg[& rX + \Pi(\mu_M - r) - C \\ & + \alpha(K\mu_A - I - \phi(I,K)) + (1-\alpha)sK - \frac{\theta_s}{2}s^2 K\Bigg]F_X \\ & + \Bigg[\frac{(\alpha\sigma_A K)^2 + 2\rho_A\sigma_A\sigma_M\alpha K\Pi + 2\sqrt{1-\rho_A^2}\sigma_A\sigma_S\alpha K\Omega}{2} \\ & + \frac{(\sigma_M^2\Pi^2) + (\Omega^2\sigma_s^2)}{2}\Bigg]F_{XX}.\end{aligned}$$

The first-order condition for consumption is simply

$$\xi e^{-\gamma C} = F_X(K,X).$$

And the optimal diversion, investment, and portfolio allocations are similarly obtained by taking the first-order conditions:

$$1 + \theta_i i = \frac{F_K}{\alpha F_X}, \quad s = \frac{1-\alpha}{\theta_s},$$

$$\Pi = -\frac{\eta}{\sigma_M}\frac{F_X}{F_{XX}} - \frac{\rho_A\sigma_A}{\sigma_M}\alpha K, \quad \Omega = -\frac{\sqrt{1-\rho_A^2}\sigma_A}{\sigma_s}\alpha K.$$

The two terms in the portfolio allocation reflect the two sources of risk hedging: the first is the standard mean–variance efficient portfolio allocation; the second is the hedge against the firm's exposure to systematic risk. The magnitude of the hedging demand for both idiosyncratic and systematic risk is scaled by αK, the relative exposure of the controlling shareholder to the firm's capital stock.

To solve the model, we conjecture and verify that the controlling shareholder's value function follows the standard form given by Merton (1971), namely,

$$F(X,K) = -\frac{1}{\gamma r}\exp\left[-\gamma r\left(X + \frac{\eta^2}{2\gamma r^2} + \frac{\zeta - r}{\gamma r^2} + \alpha G(K)\right)\right],$$

where $G(K)$ is naturally interpreted as the entrepreneur's certainty-equivalent valuation of the firm. Recall that the private benefit s_t is already given per unit of capital. We let $i_t = I_t/K_t$ denote investment per unit of capital, which implies

$$C^*(X,K) = r\left(X + \frac{\eta^2}{2\gamma r^2} + \frac{\zeta - r}{\gamma r^2} + \alpha G(K)\right).$$

Optimal investment and portfolio allocations are similarly given by

$$i^*(X,K) = \frac{G'(K) - 1}{\theta_i}, \quad \Pi^*(X,K) = \frac{1}{\sigma_M}\left(\frac{\eta}{\gamma r} - \rho_A\sigma_A\alpha K\right),$$
$$\Omega^*(X,K) = -\frac{\sqrt{1-\rho_A^2}\sigma_A}{\sigma_s}\alpha K.$$

Substituting into the HJB, after simplification we have the following ordinary differential equation (ODE) for $G(K)$:

$$rG(K) = \left(\frac{(1-\alpha)^2}{2\alpha\theta_s} + \mu_A - \rho_A\sigma_A\eta\right)K + \frac{(G'(K)-1)^2}{2\theta_i}K - \delta_K G'(K)K.$$

Controlling Shareholder's Valuation. Recall that $G(K;\alpha,\rho_A,\theta_s)$ denotes the controlling shareholder's valuation of the firm. As long as markets are complete, idiosyncratic risk can be hedged away, implying the following ODE for $G(K;\alpha,\rho_A,\theta_s)$:

$$rG(K;\alpha,\rho_A,\theta_s) = \left(\frac{(1-\alpha)^2}{2\alpha\theta_s} + \mu_A - \rho_A\sigma_A\eta\right)K + \frac{(G'(K;\alpha,\rho_A,\theta_s)-1)^2}{2\theta_i}K$$
$$- \delta_K G'(K;\alpha,\rho_A,\theta_s)K.$$

We guess the solution is in the form of $G(K) = q(\alpha)K$. After substitution we have[3]

$$q(\alpha,\rho_A,\theta_s) = 1 + (r+\delta_K)\theta_i$$
$$- \sqrt{(r+\delta_K)^2\theta_i^2 - 2\theta_i\left(\mu_A - \rho_A\sigma_A\eta - (r+\delta_K) + \frac{(1-\alpha)^2}{2\alpha\theta_s}\right)},$$

[3] The other quantities are given in the appendix.

$$i(\alpha,\rho_A,\theta_s) = (r+\delta_K) - \sqrt{(r+\delta_K)^2 - \frac{2}{\theta_i}\left(\mu_A - \rho_A\sigma_A\eta - (r+\delta_K) + \frac{(1-\alpha)^2}{2\alpha\theta_s}\right)}.$$

Outside Shareholder's Valuation. Recall that the cumulative dividends of the firm for outside shareholders follow the stochastic process

$$dY_t = (K_t\mu_A - I_t - \Phi(I_t,K_t) - s_t K_t)dt + \sigma_A\left(\rho_A dZ_t^M + \sqrt{1-\rho_A^2}dZ_t^A\right)K_t.$$

Substituting yields

$$dY_t = (\mu_A - i(\alpha) - \phi(i(\alpha,\rho_A,\theta_s)) - s(\alpha,\rho_A,\theta_s))K_t dt + \sigma_A\left(\rho_A dZ_t^M + \sqrt{1-\rho_A^2}dZ_t^A\right)K_t.$$

To determine the price of the firm, recall that we can employ risk-neutral pricing and apply a standard change of measure, which yields

$$dY_t = \big(\mu_A - \eta\rho_A\sigma_A - i(\alpha,\rho_A,\theta_s) - \phi(i(\alpha,\rho_A,\theta_s)) - s(\alpha,\rho_A,\theta_s)\big)K_t dt + \sigma_A K_t d\tilde{Z}_t^A.$$

For outside shareholders, the value of the firm is simply the present value of all future cash flows discounted under the risk-neutral measure:

$$\begin{aligned} V(K) &= \tilde{E}\left(\int_0^\infty e^{-rt}dY_t\right) = \int_0^\infty e^{-rt}\tilde{E}(dY_t) \\ &= \int_0^\infty e^{-rt}\big(\mu_A - \eta\rho_A\sigma_A - i(\alpha,\rho_A,\theta_s) - \phi(i(\alpha,\rho_A,\theta_s)) \\ &\quad - s(\alpha,\rho_A,\theta_s)\big)K_t dt. \end{aligned}$$

Recall that capital stock evolves according to the following stochastic process:

$$dK_t = (i(\alpha,\rho_A,\theta_s) - \delta_K)K_t dt.$$

We substitute this into the risk-neutral valuation formula to yield

$$V(K;\alpha,\rho_A,\theta_s) = \frac{\mu_A - \eta\rho_A\sigma_A - i(\alpha,\rho_A,\theta_s) - \phi(i(\alpha,\rho_A,\theta_s)) - s(\alpha,\rho_A,\theta_s)}{r+\delta_K - i(\alpha,\rho_A,\theta_s)}K.$$

We immediately observe that the value of the firm to the outside shareholder, $q^{\text{outsider}}K$, implies that average and marginal q^{outsider} are equal:

$$q^{\text{outsider}}(\alpha,\rho_A,\theta_s) = \frac{\mu_A - \eta\rho_A\sigma_A - i(\alpha,\rho_A,\theta_s) - \phi(i(\alpha,\rho_A,\theta_s)) - s(\alpha,\rho_A,\theta_s)}{r+\delta_K - i(\alpha,\rho_A,\theta_s)}.$$

This follows the standard formulation of the value of the firm as a ratio of the risk-adjusted expected productivity of assets after subtracting investment, adjustment costs, and private benefits to the opportunity cost of capital, i.e., risk-adjusted capital depreciation subtracting investment per unit of capital.

Price Equivalency and Welfare. Consider an investor protection parameter θ_{s1} and a correlation coefficient between the productivity shock and the tradable market risk ρ_{A1}. For a given investor protection parameter θ_{s2} we say ρ_{A2} is a price-equivalent correlation coefficient if the following equality holds:

$$q^{\text{outsider}}(\alpha,\rho_{A1},\theta_{s2}) = q^{\text{outsider}}(\alpha,\rho_{A2},\theta_{s1}).$$

In Mitts and Mansouri (2018), we discuss the properties of price equivalency. Here, I focus on a heuristic analysis of welfare implications. The following lemma holds for any increase in investor protection that also increases outside shareholder's q:[4]

LEMMA *Consider two price-equivalent allocations (θ_{s1},ρ_{A2}) and (θ_{s2},ρ_{A1}) such that $\theta_{s1} < \theta_{s2}$ and $q^{\text{outsider}}(\alpha,\rho_{A1},\theta_{s1}) < q^{\text{outsider}}(\alpha,\rho_{A1},\theta_{s2})$. The social planner strictly prefers the allocation (θ_{s1},ρ_{A2}) to (θ_{s2},ρ_{A1}), regardless of welfare weights.*

PROOF The proof is trivial. By definition of price equivalency, we must have $q^{\text{outsider}}(\alpha,\rho_{A1},\theta_{s2}) = q^{\text{outsider}}(\alpha,\rho_{A2},\theta_{s1})$. Outside shareholders are thus indifferent between the two allocations. Recall that the controlling shareholder's q is given by

$$q(\alpha,\rho_A,\theta_s) = 1+(r+\delta_K)\theta_i - \sqrt{(r+\delta_K)^2\theta_i^2 - 2\theta_i\left(\mu_A - \rho_A\sigma_A\eta - (r+\delta_K) + \frac{(1-\alpha)^2}{2\alpha\theta_s}\right)}.$$

Clearly, $\partial q(\alpha,\rho_A,\theta_s)/\partial\rho_A < 0$ for any $\rho_A \in (0,1]$. If $\theta_{s1} < \theta_{s2}$ and $q^{\text{outsider}}(\alpha,\rho_{A1},\theta_{s1}) < q^{\text{outsider}}(\alpha,\rho_{A1},\theta_{s2})$, then by price equivalency we must have $\rho_{A1} > \rho_{A2}$, so the controlling shareholder strictly prefers the allocation (θ_{s1},ρ_{A2}) to (θ_{s2},ρ_{A1}). As these are the only two agents in the economy, social welfare is higher under the allocation (θ_{s1},ρ_{A2}), regardless of welfare weights. *Q.E.D.*

Expected Return. Letting R_t^e denote the return on the firm's equity, we have by definition that dR_t^e is equal to the sum of the dividend yield dY_t/V_t and capital gains dV_t/V_t:

$$dR_t^e = \frac{dV_t}{V_t} + \frac{dY_t}{V_t}.$$

[4] A careful inspection of q^{outsider} shows that increasing θ_s may decrease q^{outsider} within a narrow parameter range where the decline in $i(\alpha,\rho_A,\theta_s)$ exceeds the decline in $s(\alpha,\rho_A,\theta_s)$. However, this does not occur for reasonable parameter values. Increasing investor protection almost always benefits outside shareholders by reducing the diversion of cash flow to private benefits.

By substitution, we can simply write

$$\begin{aligned} dR_t^e = {} & \frac{dK_t}{K_t} + \frac{(\mu_A - i(\alpha) - \phi(i(\alpha)) - s(\alpha))K_t dt}{q^{\text{outsider}} K_t} \\ & + \frac{\sigma_A(\rho_A dZ_t^M + \sqrt{1-\rho_A^2} dZ_t^A)K_t}{q^{\text{outsider}} K_t} \\ = {} & \left(i(\alpha) - \delta_K + \frac{\mu_A - i(\alpha) - \phi(i(\alpha)) - s(\alpha)}{q^{\text{outsider}}} \right) dt \\ & + \frac{\sigma_A \rho_A}{q^{\text{outsider}}} dZ_t^M + \frac{\sigma_A \sqrt{1-\rho_A^2}}{q^{\text{outsider}}} dZ_t^A. \end{aligned}$$

It is straightforward to see that the capital-asset pricing model (CAPM) holds:

$$dR_t^e = \mu^{\text{outsider}} dt + \frac{\sigma_A \rho_A}{q^{\text{outsider}}} dZ_t^M + \frac{\sigma_A \sqrt{1-\rho_A^2}}{q^{\text{outsider}}} dZ_t^A,$$

where the expected return on the firm's equity to outside shareholders is given by[5]

$$\mu^{\text{outsider}} = r + \frac{\sigma_A \rho_A}{\sigma_M q^{\text{outsider}}} (\mu_M - r),$$

i.e., the equity beta for outside shareholders is given by

$$\beta^{\text{outsider}} = \frac{\sigma_A \rho_A}{\sigma_M q^{\text{outsider}}}.$$

We can thus write the instantaneous return on the firm's equity stock as

$$dR_t^e = (r + \beta^{\text{outsider}}(\mu_M - r))dt + \frac{\sigma_A}{q^{\text{outsider}}} \left(\rho_A dZ_t^M + \sqrt{1-\rho_A^2} dZ_t^A \right),$$

which implies

$$E[dR_t^e] = r + \beta^{\text{outsider}}(\mu_M - r).$$

Under very general conditions,[6] $\partial q^{\text{outsider}}/\rho_A < 0$ and thus $\partial \beta^{\text{outsider}}/\rho_A > 0$.

Two implications follow from this analysis. First, an increase in q^{outsider} that is *not* driven by a reduction in ρ_A (i.e., by increasing θ_s) will nonetheless lead to a decline in the firm's beta. As the firm becomes more valuable, its expected return becomes less tightly linked to that of the risky asset. The firm's beta does not necessarily identify the correlation of the asset's payoffs with systematic risk. Second, recall that the controlling shareholder is better off for any decrease in ρ_A that is price-equivalent to an increase in θ_s. It thus follows that *beta-equivalence* does not imply welfare-equivalence: for any given decline in β^{outsider}, the social planner prefers that it be driven by a reduction in ρ_A rather than an increase in θ_s.

5 For a derivation see the appendix.

6 q^{outsider} is monotonically decreasing in ρ_A for any reasonable parameter value.

As shown in Mitts and Mansouri (2018), the preference for risk reduction over decreasing private benefits is unique to the complete-markets setting under CARA utility. With incomplete markets, the controlling shareholder can no longer costlessly hedge nonmarket risk and thus does not always prefer a reduction in ρ_A. Similarly, under CRRA utility, the relationship between θ_s and ρ_A is more complex due to wealth effects. Thus, the analysis here simply establishes a sufficient condition under which price equivalency does not imply welfare equivalency.

4 Discussion

This article shows that price equivalency does not imply welfare equivalency. A \$1 increase in firm value from increasing expected cash flows does not necessarily have the same welfare implications as a \$1 increase in value from reducing systematic risk exposure. But why does any of this matter? In the remaining pages, I outline three ways that the divergence between price equivalency and welfare equivalency sheds light on methodological and substantive questions in corporate and securities law.

Event Studies. The first implication is methodological. A large literature utilizes event studies to draw inferences as to the welfare effects of corporate events and policy changes. The standard interpretation is that "abnormally positive stock returns […] suggest that […] the market expects future cash flows to increase" (Bhagat and Romano, 2002, p. 410). But a price increase can reflect either higher expected cash flows or lower systematic risk exposure. And as this article has shown, welfare may be higher under one or the other for any given price increase, so separating cash-flow news from discount-rate news (Campbell and Vuolteenaho, 2004) can shed greater light on the welfare effects of a corporate event or policy change.

Thus, it may be useful to measure the effect of an event on the beta itself, i.e., the covariance of a firm's returns with those of the market and other risk factors. The asset-pricing literature studies how the beta varies over time, following theoretical work on the conditional capital-asset pricing model (Jagannathan and Wang, 1996; Lettau and Ludvigson, 2001). This can easily be done in an event study by interacting market beta with time windows or estimating the alpha across calendar-time subsamples.

The latter approach is taken by Brav, Jian, and Kim (2015), who consider the link between hedge-fund activist interventions and time-varying systematic risk. They showed that target exposure to the HML factor is statistically insignificant in the 1–3 years prior to the activist intervention, peaks in the first year after, and declines to a lower level in the following 1 to 3 years. They point out that this decline in expected returns – the "discount rate channel" – could explain some of the announcement-day returns to hedge-fund activism. And Brav et al. (2008) show substantial variation in market beta when employing calendar-time regressions of abnormal returns surrounding hedge-fund activism announcements on Schedule

13-D. Table 3 of that paper shows a beta estimate for (−12 mo, 10 mo) of 0.98, which drops to 0.39 in the (−1 mo, 1 mo) window, only to rise again to 0.92 in the (7 mo, 9 mo) window.

These findings raise the possibility that hedge-fund activists may create value by reducing the covariance of a firm's cash flows with systematic risk in addition to increasing expected cash flows. Brav et al. (2015, p. 284) note that "[a] large proportion of [names of hedge-fund activists] include words or phrases that connote value investing, such as 'value,' 'contrarian,' and 'distressed.'" To the extent that activists encourage a firm to pursue *contrarian* policies in particular, these hedge funds may create shareholder value by making a firm's cash flows more idiosyncratic. This is a slightly different conception of the role of activist investing from monitoring and reducing the consumption of private benefits of control (Kahan and Rock, 2007).

Principal Costs and Disagreement. A second, more substantive implication of this framework is that governance structures that allow for greater private benefits of control may provide offsetting benefits if they reduce systematic risk exposure. A growing literature considers how instruments like dual-class stock, which can increase agency costs, also allow founders to pursue idiosyncratic vision and thereby reduce mistakes that might arise from dispersed ownership (Goshen and Hamdani, 2016). More broadly, Goshen and Squire (2017) show a general tradeoff between agency costs and *principal costs*, such as incompetence and conflicts of interest between principals. The separation of ownership and control may reduce the latter while increasing the former, so lowering agency costs is not always value-enhancing.

The principal-cost theory in Goshen and Squire (2017) is isomorphic to the framework in this article in the following sense: idiosyncratic control facilitates a kind of contrarianism that gives rise to reduced correlation between cash flows. Disagreement between outside shareholders and a controlling shareholder implies, by definition, that the latter will choose projects that pay off in states of the world that the former believe are less likely to occur.

For a trivial example, suppose outside shareholders think it will rain but the controlling shareholder is sure it will be sunny. The controlling shareholder will choose a project that pays off when it is sunny ("sunny projects"). Outside shareholders will invest in firms with projects that pay off when it rains ("rainy projects"). This implies that the return on the controlling shareholder's projects will be less correlated with outside shareholders' portfolio returns. Suppose, for simplicity, that sunny and rainy projects have the same expected cash flow. Outside shareholders benefit when the controlling shareholder chooses sunny projects while they (in disagreement) invest in rainy projects. After all, if they are wrong, the controlling shareholder's sunny project pays off; if they are right, their portfolio of rainy projects pays off.

This example illustrates how mechanisms that allow founders to pursue an idiosyncratic vision that disagrees with outside shareholders can be value-enhancing by reducing the correlation of a firm's cash flows with outside shareholders' systematic risk exposure. Of course, the key assumption in this story is that expected cash flows are identical in the two cases: the exercise of private benefits can easily reduce the value of the firm far beyond the benefits of additional diversification. Thus, I am not suggesting that managerial entrenchment is presumptively value-enhancing, or even that it is value-enhancing a significant fraction of the time; rather, the claim is that it is *possible* that allowing founders to pursue idiosyncratic vision – even at the expense of some consumption of private benefits – may potentially benefit investors by exploiting the diversification benefits of disagreement. The empirical evidence in Brav, Jian, and Kim (2015) and Brav et al. (2008) suggests that hedge-fund activism may benefit firms either by enhancing cash flows and/or by reducing risk exposure.

Governance and Macroeconomic Stability. A large literature considers adjustment costs – frictions that impede adapting capital to more efficient uses (Hayashi, 1982). Gordon (2018) argues that adjustment costs can lead large institutional investors to resist governance changes that would destabilize labor markets or bring about political instability. Following Hansmann and Kraakman (2000), the question is "which shareholders" enter into the objective function, that is, how much weight should be given to macroeconomic stability when maximizing firm value.

This article sheds light on this question by presenting the trade-off between expected-value maximization and instability reduction in price-equivalent terms. By definition, an investor cannot diversify the beta, the firm's exposure to macroeconomic (systematic) risk. Investors with high-beta portfolios might prefer a reduction in a firm's risk exposure to an expected-value increase that has an equivalent effect on the price of the firm. This can explain why, for example, large institutional investors might resist interventions that lead to cash-flow-enhancing financial and operational change. To the extent that these changes effectively make the firm more procyclical – e.g., by hiring and firing workers and thereby exposing the firm to the ups and downs of the labor market – the firm becomes more sensitive to macroeconomic risk and thus provides less of a diversification benefit to stability-minded investors.

Of course, taken to the extreme, stability can become stagnation. Is it possible to strike a balance between the twin goals of increasing cash flow and reducing instability? Price equivalency quantifies this trade-off in terms of an observable metric (the stock price) that is routinely used to measure firm value. One might wonder if large institutional investors prefer stability at the expense of lower portfolio alpha, but this paper shows why they might prefer stability, *holding fixed* the market value of the firm. Put differently, if risk reduction yields a greater welfare increase than a price-equivalent change in expected cash flows, stability-minded corporate governance can be uniquely value-enhancing.

Appendix

Controlling Shareholder's Valuation. The other quantities are given by

$$C(X,K;\alpha,\rho_A,\theta_s) = r\left(X + \frac{\eta^2}{2\gamma r^2} + \frac{\zeta - r}{\gamma r^2} + \alpha q(\alpha)K\right),$$
$$\Pi(X,K;\alpha,\rho_A,\theta_s) = \frac{1}{\sigma_M}\left(\frac{\eta}{\gamma r} - \rho_A\sigma_A\alpha K\right),$$
$$\Omega(X,K;\alpha,\rho_A,\theta_s) = -\frac{\sqrt{1-\rho_A^2}\sigma_A}{\sigma_s}\alpha K,$$
$$s(\alpha,\rho_A,\theta_s) = \frac{1-\alpha}{\theta_s}.$$

Expected Return. To see this, simply note that

$$i(\alpha) - \delta_K + \frac{\mu_A - i(\alpha) - \phi(i(\alpha)) - s(\alpha)}{q^{\text{outsider}}} = r + \frac{\sigma_A\rho_A}{\sigma_M q^{\text{outsider}}}(\mu_M - r)$$
$$\iff i(\alpha) - \delta_K - r + \frac{\mu_A - i(\alpha) - \phi(i(\alpha)) - s(\alpha)}{q^{\text{outsider}}} = \frac{\eta\sigma_A\rho_A}{q^{\text{outsider}}}$$
$$\iff \frac{\mu_A - i(\alpha) - \phi(i(\alpha)) - s(\alpha) - \eta\sigma_A\rho_A}{q^{\text{outsider}}} = r + \delta_K - i(\alpha)$$
$$\iff q^{\text{outsider}} = \frac{\mu_A - i(\alpha) - \phi(i(\alpha)) - s(\alpha) - \eta\sigma_A\rho_A}{r + \delta_K - i(\alpha)}.$$

References

Bhagat, Sanjai, and Roberta Romano (2002), "Event Studies and the Law: Part II: Empirical Studies of Corporate Law," *American Law and Economics Review*, 4(2), 380–423.

Brav, Alon, Wei Jiang, and Hyunseob Kim (2015), "The Real Effects of Hedge Fund Activism: Productivity, Asset Allocation, and Labor Outcomes," *The Review of Financial Studies*, 28(10), 2723–2769.

—, —, Frank Partnoy, and Randall S. Thomas (2008), "The Returns to Hedge Fund Activism," *Financial Analysts Journal (FAJ)*, 64(6), 45–61.

—, —, —, and — (2015), "Hedge Fund Activism, Corporate Governance, and Firm Performance," in: William W. Bratton and Joseph A. McCahery (eds.), *Institutional Investor Activism: Hedge Funds and Private Equity, Economics and Regulation*, Oxford University Press, Oxford, pp. 261–304.

Campbell, John Y., and Tuomo Vuolteenaho (2004), "Bad Beta, Good Beta," *The American Economic Review*, 94(5), 1249–1275.

Cochrane, John H. (2005), *Asset Pricing*, revised edn., Princeton University Press, Princeton (NJ).

Gordon, Jeffrey N. (2018), "Convergence and Persistence in Corporate Law and Governance," in: Jeffrey N. Gordon and Wolf-Georg Ringe (eds.), *The Oxford Handbook of Corporate Law and Governance*, Oxford University Press, Oxford, pp. 28–55.

Goshen, Zohar, and Assaf Hamdani (2016), "Corporate Control and Idiosyncratic Vision," *The Yale Law Journal*, 125(3), 560–795.

— and Richard Squire (2017), “Principal Costs: A New Theory for Corporate Law and Governance,” *Columbia Law Review*, 117(3), 767–829.

Hansmann, Henry, and Reinier Kraakman (2000), “The Essential Role of Organizational Law,” *The Yale Law Journal*, 110(3), 387–440.

Hayashi, Fumio (1982), “Tobin’s Marginal q and Average q: A Neoclassical Interpretation,” *Econometrica*, 50(1), 213–224.

Hayek, F. A. (1945), “The Use of Knowledge in Society,” *The American Economic Review*, 35(4), 519–530.

Jagannathan, Ravi, and Zhenyu Wang (1996), “The Conditional CAPM and the Cross-Section of Expected Returns,” *The Journal of Finance*, 51(1), 3–53.

Jensen, Michael C., and William H. Meckling (1976), “Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure,” *Journal of Financial Economics*, 3(4), 305–360.

Johnson, Simon, Rafael La Porta, Florencio Lopez-de-Silanes, and Andrei Shleifer (2000), “Tunneling,” *The American Economic Review*, Papers and Proceedings, 90(2), 22–27.

Kahan, Marcel, and Edward B. Rock (2007), “Hedge Funds in Corporate Governance and Corporate Control,” *University of Pennsylvania Law Review*, 155(5), 1021–1093.

La Porta, Rafael, Florencio Lopez-de-Silanes, Andrei Shleifer, and Robert Vishny (2002), “Investor Protection and Corporate Valuation,” *The Journal of Finance*, 57(3), 1147–1170.

Lan, Yingcong, Neng Wang, and Jinqiang Yang (2012), “Investor Protection, Diversification, Investment, and Tobin’s q,” September 19, https://www0.gsb.columbia.edu/mygsb/faculty/research/pubfiles/1749/InvestorProtection_Lan_Wang_Yang_Sept2012.pdf, accessed November 3, 2018.

LeRoy, Stephen F., and Jan Werner (2014), *Principles of Financial Economics*, 2nd edn., Cambridge University Press, New York.

Lettau, Martin, and Sydney Ludvigson (2001), “Resurrecting the (C)CAPM: A Cross-Sectional Test when Risk Premia Are Time-Varying,” *Journal of Political Economy*, 109(6), 1238–1287.

Merton, Robert C. (1973), “An Intertemporal Capital Asset Pricing Model,” *Econometrica*, 41(5), 867–887.

Mitts, J., and S. Mansouri (2018), “Disagreement and the Value of the Firm,” unpublished Manuscript, Columbia Law School, New York.

Stulz, René M. (2005), “The Limits of Financial Globalization,” *The Journal of Finance*, 60(4), 1595–1638.

Weil, Philippe (1992), “Equilibrium Asset Prices with Undiversifiable Labor Income Risk,” *Journal of Economic Dynamics & Control*, 16(3–4), 769–790.

Joshua Mitts
Columbia Law School
435 West 116th Street
New York, NY 10027
USA
joshua.mitts@law.columbia.edu