RAND Journal of Economics
Vol. 38, No. 3, 2007
pp. 626–647

Do investors forecast fat firms? Evidence from the gold-mining industry

Severin Borenstein*

and

Joseph Farrell**

Conventional economic theory assumes that firms minimize costs given output, but news articles and managers indicate that firms cut costs when they are in economic distress and grow fat when they are relatively wealthy. Under conventional theory, firm value is convex in the price of a competitively supplied input or output, but we find that the stock values of many gold-mining companies are concave in the price of gold. We show that this is consistent with fat accumulation when a firm grows wealthy. We then address alternative explanations and discuss where fat in these companies might reside.

1. Introduction

■ Organizations do not always minimize costs or maximize value. There can be sheer inefficiency or rent dissipation. In this article, we take a simple empirical approach to such "fat" by testing a rather general theoretical property of value maximization. The empirical results suggest that many gold-mining companies grow fat when they get rich and that the amounts concerned may be quite large.

* University of California, Berkeley, and NBER; borenste@haas.berkeley.edu.
** University of California, Berkeley; farrell@econ.berkeley.edu.
For research support, we are grateful to the Alfred P. Sloan Foundation, via a grant to the Industrial Technology and Productivity Project of the National Bureau of Economic Research. For helpful comments, we thank Alan Auerbach, Judith Chevalier, Erwin Diewert, Daniel McFadden, Kevin M. Murphy, Jim Poterba, David Scharfstein, and Peter Tufano, and participants in numerous seminars and conferences. We also thank two anonymous referees and the editor, Rob Porter, for their valuable input. T.T. Yang, David O'Neill, and Evan Rose provided outstanding research assistance. We also thank gold-mining managers and government personnel who helped us understand industry issues: Earl Amey (U.S. Geological Survey), Paul Bateman (The Gold Institute), Margo Bergeson (Alta Gold), Pierre-Paul Bleau (Cambior), Mike Brown (Amax Gold), Ronald Cambre (Newmont Mining), Jerry Cooper (Asarco), Steve Dawson (Campbell Resources), Larry Drew (Hecla), Jane Engert (U.S. EPA), Tom Ferrell (Newmont Mining), Robert Gilmore (Dakota Mining), Shannon Hill (Placer Dome), Denise Jones (California Mining Association), Robbin Lee (Echo Bay), John Lutley (The Gold Institute), Greg Pelka (California State Lands), Michael Steeves (Homestake Mining), Les Van Dyke (Battle Mountain Gold), Dennis Wheeler (Coeur D'Alene), Richard Young (Barrick Gold), and Les Youngs (California Office of Mine Reclamation). Any remaining errors are the sole responsibility of the authors.

We begin by recalling why simple maximizing theory predicts that a firm's value should be a convex function of exogenous prices. The result is a simple application of real-option theory; it requires no assumptions about functional forms or elasticities of the industry demand function or the firm cost function. We then show that the convexity result can fail if the firm grows "fatter"—that is, dissipates a larger share of value—as the firm becomes richer.

We test the theory in the gold-mining industry. The industry is attractive for this test because gold prices are quite exogenous to a gold-mining firm and cause large changes in the value of the firm. We estimate the relationship between the price of gold and the stock market equity valuation of 17 gold-mining firms. We find that in nearly half these firms, the relationship is significantly *concave*. We address potential alternative explanations for this surprising finding and argue that they are unlikely to explain as much concavity as we find.

We believe that the approach we present convincingly shows that some form of fat exists and is important. Unfortunately, however, this approach only *bounds* from below the quantity of fat; it does not permit direct *estimation* of the quantity. As a result, it does not allow for useful cross-firm comparisons of fat.

Section 2 recalls why economic theory predicts that the value of a fat-free firm is convex in exogenous prices. We then show how fat could disrupt this result, and how, if concavity of the value function is due to fat, its magnitude implies a lower bound on fat. Section 3 applies the theory empirically to the gold-mining industry by estimating the response of the equity value of gold-mining firms to changes in the price of gold. Section 4 argues that various other potential explanations seem unlikely to account for the significant concavity that we find. Section 5 asks where such fat might be. Section 6 concludes.

2. How firm value responds to price shocks

■ We begin by presenting the standard theoretical argument that the value of a fat-free firm is convex in prices. We then explore the relationship if a firm gets "fat" as it grows richer.

□ **The value response of a fat-free firm.** The maximized value of a firm is a (nonstrictly) *convex* function of any exogenously determined price it faces, holding constant other prices and the constraints and terms of trade facing the firm. This fundamental (and well-known) result holds whether the price is that of an input, an output, or a good that is sometimes an input and sometimes an output. The result does not depend on assumptions about production technology (beyond assuming that prices do not affect technological possibilities) or about the shape or elasticity of demand.

To recall why, note that for any fixed production plan, the firm's value is linear in each price.[1] For example, if a gold-mining firm ignored changes in the price of gold and just mined a given quantity, x ounces, the firm's value would be $v(p_g, x) \equiv p_g x - C(x)$, which is linear in the price of gold, p_g. If p_g rose from \$300 to \$400 per ounce, $v(p_g, x)$ would rise by \$100x, just as it would if p_g rose from \$400 to \$500 per ounce. Thus, given a production plan with output x, the firm's value $v(p_g, x)$ would be linear in the gold price p_g.[2]

However, the firm can change production plan, and typically will do so in maximizing value. Gold-mining firms vary output with the price of gold: Moel and Tufano (2002) find that firms often close and reopen gold mines in response to changes in gold prices, consistent with real-option analysis. Because the firm can profitably change production plans, its maximized value, as a function of price, is an upper envelope of straight lines, $V(p_g) \equiv \max_x v(p_g, x)$, and hence is convex, as shown in Figure 1.

As the argument suggests, by the envelope theorem $V'(p) = x(p)$, so the slope of the value function should be equal to the firm's anticipated output quantity x, or, in the long-run valuation

[1] In this article, we call a function "linear" if it has constant slope, whether or not it passes through the origin.

[2] A similarly linear relationship would hold if it were an input such as labor, rather than the firm's output, whose price was changing while the firm kept the same production plan.

FIGURE 1

UPPER ENVELOPE OF STRAIGHT LINE VALUE FUNCTIONS IS CONVEX

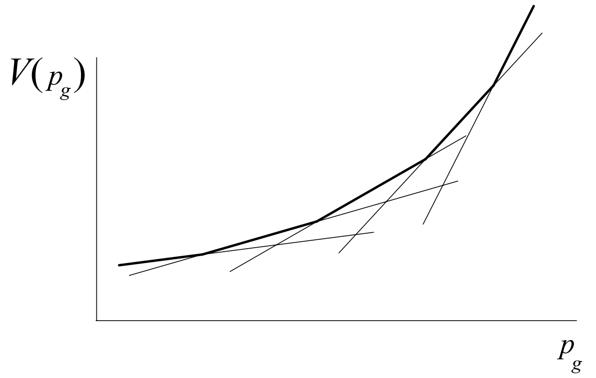

of a gold-mining firm, something like the firm's economic reserves.[3] More interestingly, the convexity of V is closely related to the "real-option" value to the firm of being able to adjust quantities in response to price changes. Because $V'(p) = x(p)$, it follows that $V''(p) = x'(p)$, so a natural measure of the curvature of the V function, $pV''(p)/V'(p)$, the elasticity of the slope of V with respect to price, is equal to the elasticity of the firm's supply with respect to price, $p\,x'(p)/x(p)$.

Although this proof of convexity of the firm's value function is straightforward when only one price varies, other relevant prices may also vary. For example, when the spot price of gold changes, expected future prices of gold presumably also change. When multiple prices simultaneously change, the theoretical convexity result has a natural generalization: maximized value V is now an upper envelope of hyperplanes rather than of straight lines. This implies the well known[4]

Proposition 1. Consider a firm that maximizes value V taking as given (input and output) prices $\mathbf{p} \equiv (p_1, \ldots, p_N)$. The maximized value $V(\mathbf{p})$ is a convex function of the vector $\mathbf{p}$.

If we observed the entire vector $\mathbf{p}$ of relevant prices, this proposition would let us test directly for value maximization. Also, prices that do not change can of course be dropped from the price vector. However, some omitted prices may well vary in the sample. Indeed, what we do empirically below is track the empirical relationship between *one* price—"the" price, p_g, of gold, the primary output—and the stock market's assessed equity value of the firm. How is this relationship affected if other, excluded, prices change in a way that is correlated (in the sample) with p_g? We address two versions of this question. First, we consider prices that the firm and investors can observe, but that we omit from our regressions. Second, we consider future prices that are uncertain at the observation date.

Omitted (known) prices. The effect of omitting a relevant price that is observed by market participants at the observation date differs according to whether it is linearly or nonlinearly related to the included price(s).

Suppose first that certain excluded prices are linearly related to p_g in the sample. Then the price vectors in the sample lie on a straight line in price space, and the observed function $\hat{V}(p_g)$

[3] Strictly, the slope of the value function is the discounted value of the anticipated movement in future prices times outputs at those future dates, ignoring hedging.

[4] For an accessible statement see, for example, Kreps (1990). The earliest published source we have found is Gorman (1968); see also McFadden (1966). Daniel McFadden (personal communication, 2006) reports that the result goes back to the nineteenth century.

is the slice of the convex value function $V(\mathbf{p})$ that lies above that straight line. Consequently, $\hat{V}$ is convex along that straight line, and empirically will appear convex as an apparent or reduced-form function of p_g alone.

In particular, of course, a gold-producing firm's maximized value will depend on the future prices of gold. As just noted, if those prices were deterministically and linearly related to our single spot gold price measure, p_g, convexity would still hold. Similarly, if the price of an input, such as skilled labor, changed linearly with the price of gold, theory would still predict a convex estimated value function, even if we omitted wage rates from our regressions.

The effect is a bit less clear if the relationship between omitted and included prices is nonlinear. Suppose for simplicity that just two prices affect V, and examine the observed relationship between p_1 and V when (i) V is indeed a convex function of the full price vector (p_1, p_2), and (ii) $p_2 = f(p_1)$, where f is nonlinear (but note that any causality between p_1 and p_2 is not important here). Then the reduced-form or observed relationship between V and p_1 will be $\hat{V}(p_1) \equiv V(p_1, f(p_1))$. To study the convexity of $\hat{V}$, we calculate:

$$\hat{V}''(p) = V_{11} + 2f'(p)V_{12} + f''(p)V_2 + [f'(p)]^2 V_{22}. \tag{1}$$

Thus the observed relationship $\hat{V}$ will be convex unless

$$f''(p)V_2 < -\left[V_{11} + 2f'(p)V_{12} + (f'(p))^2 V_{22}\right], \tag{2}$$

where the expression in square brackets is positive by convexity of V in the vector (p_1, p_2). Thus, $\hat{V}$ will still be convex unless (i) $f(\cdot)$ is "sufficiently" nonlinear; (ii) p_2 is "sufficiently" important in V; and either (iiia) f is convex and good 2 is an input or else (iiib) f is concave and good 2 is an output. Of course, when f is linear, we recover the result of the discussion above.

Uncertain (future) prices. If some prices are uncertain and their distribution is unaffected by changes in the observed price p_1, then convexity follows because value can be expressed as an expected value:

$$\hat{V}(p_1) = \int V(p_1, p_2)\,\mathrm{d}F(p_2 \mid p_1), \tag{3}$$

and this is a sum of convex functions of p_1 provided that the distribution function F does not shift with changes in the observed price p_1. This argument holds however much or little the firm will be able to reoptimize as further information about currently unobserved prices arrives: that affects the shape of V as a function of p_2, but that is irrelevant for this argument.

Harder questions arise if the distribution of unobserved prices varies with the observed price, as is likely, especially when the unobserved prices are future spot prices of gold. When the expected value of each future price moves linearly in the observed spot price, so that, for instance, $\mathrm{E}\,(p_{t+1} \mid p_t) = a + bp_t$, there is a natural intuition that convexity will carry over. V is convex in future prices as well as in today's price, and if expected future prices are linear in today's price, one might expect $\hat{V}$ to be convex in today's price. This argument would be just a special case of the analysis in equation (2) if the relationship among prices were deterministic. It also goes through, by the previous discussion, if the firm cannot respond to later news about future prices, because firm value is then linear in these prices. If, however, the firm will be able to respond to future prices, the option value resulting from the variability in those prices may vary with today's price. For example, imagine that extreme values of p_1 (high or low) correspond to low conditional variances of future prices, so the conditional variance is an inverted U-shape as a function of p. If the option value is an important part of expected profits in the second period, extreme values of p_1 would then correspond to low expected second-period profits, and potentially to low present values.

To investigate this problem, consider an illustrative two-period model. At the beginning of period 1, the firm has a stock S of ore. It learns the first-period price, p_1, and then chooses first-period extraction (and sales), x_1. Its costs in the first period of extracting x_1 are $x_1^2/(2S)$, so its marginal cost is increasing linearly in x_1 and decreasing in total stock or reserves. It chooses x_1 knowing the conditional distribution of the second-period price p_2, and it knows that at the beginning of period 2, it will learn p_2 and will then choose second-period output x_2, at cost $x_2^2/(2[S - x_1])$.

Given x_1 and p_2, x_2 maximizes $p_2 x_2 - x_2^2/(2[S - x_1])$, so second-period profits are $(S - x_1)p_2^2/2$. Consequently, given p_1, x_1 maximizes

$$\tilde{V}(p_1, x_1) \equiv p_1 x_1 - x_1^2/(2S) + \frac{\delta}{2}(S - x_1)\mathrm{E}\big[p_2^2 \mid p_1\big], \tag{4}$$

where δ is the discount factor.

The reduced-form value function $\hat{V}(p_1)$ is of course simply $\max_{x_1} \tilde{V}(p_1, x_1)$. By the envelope theorem, $\hat{V}'(p_1) = \partial \tilde{V}/\partial p_1$, so differentiating again,

$$\hat{V}''(p_1) = x_1'(p_1)\frac{\partial^2 \tilde{V}}{\partial p_1 \partial x_1} + \frac{\delta}{2}(S - x_1)\frac{\mathrm{d}^2}{\mathrm{d}p_1^2}\mathrm{E}\big[p_2^2 \mid p_1\big]. \tag{5}$$

From the implicit-function theorem, $x_1'(p_1)$ has the same sign as the mixed partial derivative of $\tilde{V}$. Consequently, $\hat{V}$ is convex unless $\mathrm{E}[p_2^2 \mid p_1]$ is sufficiently concave in p_1; and $\mathrm{E}[p_2^2 \mid p_1] = (\mathrm{E}[p_2 \mid p_1])^2 + \mathrm{var}\,[p_2 \mid p_1]$, and the first term is convex in p_1 if price follows a martingale. Thus, the observed value function will be convex unless the conditional variance is quite concave.

In Section 4 below, we explore how the conditional variance changes with gold prices and, as casual observation might suggest, find no evidence of such concavity of the conditional variance. If anything, extremely high values of p_1 correspond to high, rather than low, conditional variances.

□ **Value response and fat accumulation.** Although value functions should be convex in price if firms maximize value, Leibenstein (1966) and later work on agency and free cash flow (e.g., Jensen 1986) suggest that firms accumulate fat when they become wealthy and financial constraints loosen. If fat is an increasing and convex function of wealth, it could make the firm's net-of-fat value, $V - F$, concave.

Consider a gold-mining company that would have value V if it operated with no fat. Its actual value will be $S = V(p_g) - F(V(p_g))$, where $F(V)$ is the present value of fat (profit dissipated through inefficiency), which we take to be a function of V. Differentiating with respect to p_g,

$$S'(p_g) = V'(p_g)[1 - F'(V(p_g))], \tag{6}$$

which will have the same sign (presumably positive) as $V'(p_g)$ if $F'(V) < 1$, that is, fat does not consume more than 100% of marginal wealth changes. Differentiating (6) with respect to p_g gives

$$S''(p_g) = V''(p_g)[1 - F'(V(p_g))] - V'(p_g)^2 F''(V(p_g)). \tag{7}$$

Hence,

$$\frac{S''(p_g)}{S'(p_g)} = \frac{V''(p_g)}{V'(p_g)} - \frac{F''(V)}{1 - F'(V(p_g))}V'(p_g). \tag{8}$$

The first term on the right in (8) is the (nonnegative) ratio of marginally economic reserves (those barely worth extracting at price p_g) to total economic reserves (all those worth extracting at price p_g). The second term is a measure of the convexity of the fat function. If the fat function is sufficiently convex, relative to the firm's real-option opportunity to reoptimize when p_g changes (the first term in (8)), then $\frac{S''(p_g)}{S'(p_g)}$ will be negative, making S concave in price.

Therefore, if empirically S is concave in p_g, rejecting simple versions of full maximization, this could suggest (at least investor expectations of) fat. It can also imply a lower bound on expected fat. Because

$$F(V(p_g)) \equiv V(p_g) - S(p_g), \tag{9}$$

FIGURE 2

CONSTRUCTING A LOWER BOUND ON FAT FROM VALUE FUNCTION CONCAVITY

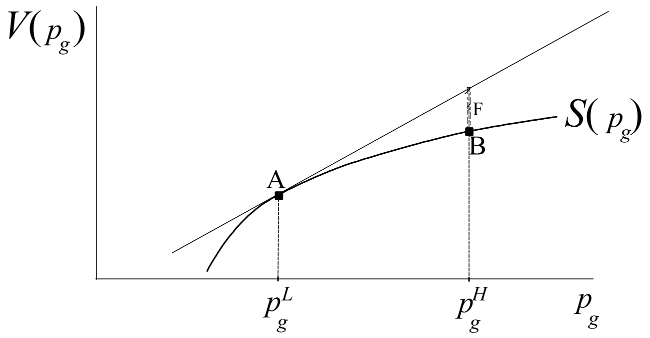

we can differentiate and divide by $V'(p_g)$ to get

$$F'(V(p_g)) = 1 - \frac{S'(p_g)}{V'(p_g)}. \tag{10}$$

Now consider observations at two prices: a low price, p_g^L, and a higher price, p_g^H, as in Figure 2. As theory tells us that $V'(p_g^H) \geq V'(p_g^L)$, and as we presume that fat increases in wealth and hence in price (i.e., F' and V' are positive), we have $V'(p_g^H) \geq V'(p_g^L) \geq S'(p_g^L)$, whence

$$F'\left(V\left(p_g^H\right)\right) = 1 - \frac{S'\left(p_g^H\right)}{V'\left(p_g^H\right)} \geq 1 - \frac{S'\left(p_g^H\right)}{S'\left(p_g^L\right)}. \tag{11}$$

This gives us an observable lower bound on the fraction of the marginal dollar of wealth gain from an increase in p_g near p_g^H that is dissipated as fat, that is, the quantity $F'(V(p_g^H))$, or "marginal fat." It is one minus the slope on the S function at point B divided by the slope at point A. The bound is strictly positive when S is concave so that $S'(p_g^H) < S'(p_g^L)$.

We also get an observable lower bound on the total rent dissipation. From convexity of V,

$$V\left(p_g^H\right) \geq V\left(p_g^L\right) + \left(p_g^H - p_g^L\right)V'\left(p_g^L\right). \tag{12}$$

As fat is non-negative (so $V \geq S$) and, we assume, weakly increasing in wealth (so $V' \geq S'$), the right-hand side is at least equal to

$$S\left(p_g^L\right) + \left(p_g^H - p_g^L\right)S'\left(p_g^L\right), \tag{13}$$

so

$$F\left(V\left(p_g^H\right)\right) \equiv V\left(p_g^H\right) - S\left(p_g^H\right) \geq \left(p_g^H - p_g^L\right)S'\left(p_g^L\right) - \left[S\left(p_g^H\right) - S\left(p_g^L\right)\right], \tag{14}$$

and this lower bound on total fat at p_g^H is positive when S is concave.

Equations (11) and (14) let us infer bounds on marginal and total fat. Because we do not observe V', we cannot *estimate* fat but only bound it, on the assumption that fat causes the concavity of $S(\cdot)$. The lower bounds will underestimate fat if, as one would expect, $V(\cdot)$ is *strictly* convex or if the firm has some fat even at p_g^L (that is, $F(p_g^L) > 0$ and $F'(p_g^L) > 0$). Of course, these supposed lower bounds will overestimate fat if concavity is caused by other factors such as those we discuss in Section 4.

This theory of fat accumulation is obviously related to a free cash flow view of managerial behavior, but there are important differences. First, whereas empirical work on free cash flow relies on accounting measures, our approach works through investors' responses to an exogenous shock to the firm's wealth. By incorporating investors' expectations, our approach may better capture changes in firm behavior that do not quickly affect cash flow. For instance, managers might respond to an increase in the price of gold by committing to higher wages, establishing attractive pension plans, or planning new capital investments, whose cash-flow cost may predictably arrive years later. Thus, our approach incorporates information that is missing, or mistimed, in free cash flow analysis.

3. Gold prices and the valuation of gold-mining companies

■ We study the effects of wealth changes on corporate fat in the gold-mining industry, because the frequent shocks to the price of gold are exogenous to the gold-mining companies we examine and translate directly into wealth shocks for those firms.[5]

Gold-mining companies are almost certainly price takers in the gold market. The market for gold is worldwide, because of the metal's high value-to-weight ratio and homogeneity, and no producer controls a large share of the annual extraction of new gold.[6] More importantly, final demand for gold can be fulfilled from existing stock. Annual production of gold from mines worldwide is about 2% of existing stock of the metal.[7] Thus, unilateral market power appears to be absent. Coordinated oligopoly interactions seem extremely unlikely, given the large number of diverse gold-mining firms and other holders of gold. Nor do theories of raising rivals' costs apply here.

For the analysis of Section 2 to apply directly, changes in the price of the important input or output should be exogenous to the firms observed. This means not only that no firm has market power but also that price movements are not driven by aggregate shocks to the observed firms, such as new gold discoveries by them or revisions in their estimated reserves.[8] In fact, gold price changes are almost uniformly the result of demand-side news: world events that change the attractiveness of gold as a store of wealth, trends in the demand for gold jewelry, or policy decisions of central banks to hold more or less gold.[9] We searched the *Wall Street Journal* over the 28-year span of our sample for articles about gold prices and found almost no mention of gold supply (from gold mines) as a cause of gold price changes. Discussions with investor relations personnel at several gold-mining companies also failed to uncover cases in which supply shocks from mines were thought to have significantly affected prices. Even the fraudulent Bre–X incident in May 1997, in which an area of Indonesia that had been touted as the largest gold find in history turned out to have no economic supplies, did not significantly affect the price of gold.[10]

The value of a gold-mining firm should depend on spot and all information about future prices of gold, but we analyze the relationship empirically using one (near-term futures) price. We subject this assumption to robustness tests below, we believe it is sensible because the prices move together very closely. Gold is actively and thickly traded and can be stored cheaply (relative to its value), so arbitrage would be comparatively easy if traders detected systematic departures from a martingale. Indeed, augmented Dickey-Fuller tests using our weekly gold price series for 1977–2004 fail to reject a unit root.[11] This is consistent with the findings of Pindyck (1993). Selvanathan (1991) found that a random-walk hypothesis performed better than a panel of gold price forecasters. Figure 3 shows the price of gold over our sample period (in constant 2004 dollars), and Table 1 gives descriptive statistics.

[5] A number of papers examine the related issue of the relationship between competition and firm productivity. See Schmitz (2005) and references therein.

[6] The largest producer in 2004, Newmont Mining, mined about 9% of world extraction and held about 7% of world reserves.

[7] According to U.S. Geological Survey data, existing stock is 128,000 tons and annual production is about 2500 tons. See minerals.usgs.gov/minerals/pubs/mcs/2006/mcs2006.pdf.

[8] In Borenstein and Farrell (1996), we analyzed the value response of oil companies to changes in the price of oil, but abandoned that analysis because the oil market lacks these advantages.

[9] One could, of course, regard these central banks' decisions as supply shocks, but they are not shocks to the supply of the firms we study.

[10] On May 6, while the stock of Bre–X fell 97% in value (confirming that the lack of economic supplies was news to the market), the price of gold *fell* about \$2/oz.

[11] The test statistic is -2.09 and the 95% critical value is -3.12.

FIGURE 3

REAL GOLD PRICES, 1977–2004

TABLE 1 Descriptive Statistics on Gold Price, 1977–2004

		Percentiles	
		(Constant 2004 Dollars)	
Mean	569.27	10th	325.17
Std dev	244.20	25th	409.15
Min	278.32	50th	501.55
Max	2013.22	75th	670.82
		90th	821.39

To analyze the effect of gold prices on a gold-mining firm's stock market value, one wants to control for market-wide stock price movements, which may represent, among other things, interest-rate changes or expected changes that would affect gold mine stock prices directly.[12] Thus, we begin with the standard capital assessment pricing model (CAPM) market model of equity returns:

$$R_{it} = R_{ft} + \beta_i(R_{mt} - R_{ft}) + \epsilon, \tag{15}$$

where R is the rate of return, the i subscript refers to the observed firm, the m subscript refers to the market, and the f subscript refers to the risk-free rate of return. We multiply both sides of (15) by the stock value of the firm at $t - 1$ to get the equation in terms of the change in firm value:[13]

$$R_{it}S_{it-1} = \Delta S_{it} = R_{ft}S_{it-1} + \beta_i(R_{mt} - R_{ft})S_{it-1} + \epsilon S_{it-1}, \tag{16}$$

[12] Over our sample period, the correlation between the return on the market index and the return on gold futures is about 0.05, which is significantly different from 0 at the 5% level.

[13] We focus on just the equity value of the firm. As we discuss in Section 4, this should strengthen the predicted convexity.

where Δ indicates the difference between the period t and period $t-1$ value of the variable. We then specify explicitly the effect of the price of gold, which would otherwise be included in the error term.

Recall that we are interested in the curvature of the relationship between S and p_g. This might be measured by the second derivative of a levels equation. Because our equation is in differences, we include the difference/derivative of a quadratic relationship between stock value and the price of gold. That is, if $S = \gamma_0 + \gamma_1 p_g + \gamma_2 p_g^2$, then $\mathrm{d}S = \gamma_1 \mathrm{d}p_g + 2\gamma_2 p_g \mathrm{d}p_g$. So, we estimate the equation

$$\Delta S_{it} - R_{ft} S_{it-1} = \alpha_1 \Delta p_{gt} + \alpha_2 p_{gt-1} \Delta p_{gt} + \alpha_3 S_{it-1}(R_{mt} - R_{ft}) + S_{it-1}\epsilon, \tag{17}$$

where S is the stock market value of the firm, p_g is the price of gold, R_{mt} is the return on a value-weighted stock market index,[14] and αs are parameters. In this model, α_3 is the estimate of the CAPM β.[15] The coefficient α_2 indicates the convexity (if $\alpha_2 > 0$) or concavity (if $\alpha_2 < 0$) of the relationship between the price of gold and the value of the firm.

We examine the stock market values of 17 gold-mining companies that are traded in the United States or Canada. We arrived at this dataset by examining lists of U.S. and Canadian gold producers and including each firm that (i) produced at least 10,000 ounces of gold in 1996, (ii) mined gold predominantly or exclusively in the United State, Canada, and Australia (we used this criterion to minimize the effect of political risk), (iii) was primarily in the gold-mining business, and (iv) was publicly traded and is covered by Center for Research in Security Prices (CRSP) stock market data. This produced 21 firms. We then eliminated 4 firms for which fewer than 104 weekly stock observations (2 years of observations) were available. For all 17 firms used in the analysis, estimation of equation (17) with just a linear gold price term indicated that the value of the firm has a positive and statistically significant relationship to the price of gold.

Our full sample period is weekly observations for January 1977 through December 2004, a total of 1458 weeks.[16] No firm is in the sample for the entire period, however. Some firms came into existence after 1977, whereas others were delisted, and ultimately ceased to operate, prior to 2004. Some firms have recently diversified and gold mining has become a relatively small share of their operations, so we drop recent years of operations for these firms. A few firms also made major purchases of other gold-mining companies in the late 1990s and early 2000s.

The stock market values are taken from CRSP data.[17] We use the nearest-contract gold futures price (traded on the COMEX division of the New York Mercantile Exchange) to represent the price of gold. Although that contract changes every other month, our Δp_{gt} variable is always the change for a given contract, not a comparison of prices on two different contracts. For the risk-free rate, we use the 1-year T-bill yield on the day of observation transformed to a weekly interest rate. For each company, we use weekly observations (closing price on the last trading day of each week) to estimate the value of the firm as a function of the price of gold.

Because we expect rent dissipation or fat to depend on the firm's real wealth, we deflate all variables. We deflate S, p_g, and the market index on which R_m is based to 2004 dollars using the consumer price index (CPI) (all items—urban consumers). We translate the nominal T-bill yield used for R_f to a real yield by $R_f^r = \frac{1+R_f}{1+\pi} - 1$, where π is the inflation rate calculated from the CPI for the month of the observation.

[14] Using an equally weighted index instead changes the results minimally.

[15] Inclusion of a constant term in (17) yields practically identical results. Similarly, rather than subtracting $R_{ft} S_{it-1}$ from both sides of (16)—implicitly assuming its coefficient is 1—inclusion of this term on the right-hand side makes virtually no difference in the results.

[16] We drop from the sample the weeks of September 10, 17, and 24 because of the September 11, 2001, terrorist attacks, which disrupted financial markets.

[17] We include weeks in which the stock goes ex-dividend or the number of shares outstanding changes, but adjust firm valuation for these changes. Dropping these weeks results in the loss of about 10% of all observations and has virtually no effect on the results.

The error term in the regression we estimate may be heteroskedastic, both because the equation is in terms of the value of the firm (as indicated in (16)) which changes over time, and because exogenous factors affect the volatility of stock market returns. We address this problem by estimating the regression using GLS, explicitly controlling for heteroskedasticity caused by the presence of S_{t-1} in the error term. We do the GLS estimation by dividing both sides of equation (17) by S_{t-1}. We then report White heteroskedastic-consistent standard errors to control for other heteroskedasticity. We have also carried out the analysis without the GLS correction, but just implementing the White correction to the standard errors, with very similar results.[18]

□ **Estimation of a quadratic value function.** We begin by separately estimating (17) for each of the 17 firms by GLS. The sample periods differ across firms, but each regression includes at least 287 observations and the median number of observations is 911. The results are shown in Table 2. The estimated second derivative is negative for 11 of the 17 firms, and significantly negative (at the 5% level) for 8 of them. Of the 6 estimated positive second derivatives, only 1 is statistically significant. The z-statistic for the 17 estimated second derivatives is -27.54 with a standard error of $\sqrt{17} = 4.12$, which is significant at the 1% level. Thus, we find a *concave* relationship between the price of gold and the values of many of the gold-mining firms.

The other parameters estimated appear reasonable. The implied first derivative of stock market value with respect to the price of gold, $\alpha_1 + \alpha_2 \cdot p_g$, is positive for each firm at the median price of gold in the sample and for nearly all gold price values that occur while the firm is in the sample. The CAPM β parameter estimated for these firms varies, but is significantly below 1 in all cases.[19]

To interpret the magnitude of the curvature of the estimated value function, we create a benchmark slope for each firm in its lean state. We calculate the estimated slope of each firm's value function when the price of gold is \$409.15, its 25th percentile value in the full 28-year sample. The first column of Table 3 shows the estimated change in slope when the price of gold increases to \$501.55, its median value in the full sample.

Graphically, this calculation compares the slope at point B to the slope at point A in Figure 2. In terms of our equation (11), this is $S'(p_g^H)/S'(p_g^L) - 1$. For Campbell Resources, for example, if the statistically significant 5.6% estimated decline in the slope were precise and were due solely to fat, this would suggest that when the price of gold increases slightly starting from its median level, at least 5.6% of the incremental gain is dissipated, that is, is not reflected in increased shareholder wealth. Recall that, taking the point estimate as correct and assuming all concavity is due to fat, this is a lower bound, as the V function is (weakly) convex.

The z-statistic discussed above is one way to aggregate our data across firms. Another is to study the response of a portfolio of gold-mining firms to changes in p_g. We do this by taking a weighted average of the slopes of the value functions, with each firm's weight being its average market capitalization while it is in the sample. We then again calculate how the slope of S is estimated to change if p_g rises from its 25th percentile to the 50th percentile value, as a percentage of the slope when p_g is at its 25th percentile. The result is an estimated decline of 11.5%, and is significant at the 5% level.[20]

If this concavity were due solely to fat, our estimates would also imply a lower bound on the *total* fat that accumulates when p_g increases, as a proportion of the theoretical increase

[18] There is also a potential econometric issue of cointegration of firm value and the (nearest-contract) futures price of gold. We tested for cointegration and found that in all cases, the null hypothesis of no cointegration was strongly rejected: t-statistics above 18 in all cases compared to a critical value that varies with number of observations, but is around 3.

[19] Using an equally weighted, rather than value-weighted, market index moves the β estimates for all firms closer to 1, but otherwise has minimal effects on the results. We also estimated the βs without the gold price terms and found very similar estimates.

[20] The variance of this estimate is calculated on the assumption that the estimates for each firm are statistically independent.

TABLE 2 Results from Quadratic Value Function Regressions

Company	Δp_{gt}	$p_{gt-1}\Delta p_{gt}$	CAPM β	R^2	Obs	Period
Alta	122	0.00	0.46	0.153	1162	1977–1999
	(33)	(0.04)	(0.13)			
Amax	1484	5.07	0.42	0.353	564	1987–1998
	(1557)	(2.96)	(0.18)			
Barrick	−918	2.61	0.63	0.110	517	1985–1994
	(1144)	(1.94)	(0.19)			
Campbell$^-$	136	−0.07	0.59	0.098	1295	1977–2001
	(18)	(0.01)	(0.14)			
Canyon$^-$	392	−0.54	−0.30	0.194	983	1986–2004
	(136)	(0.19)	(0.19)			
Coeur D'Alene$^-$	399	−0.21	0.52	0.302	596	1979–1990
	(54)	(0.05)	(0.14)			
Dakota	−85	0.45	0.07	0.076	428	1989–1998
	(216)	(0.42)	(0.39)			
Echo Bay	1131	1.64	0.02	0.230	1004	1983–2003
	(440)	(0.91)	(0.14)			
Getchell$^-$	1948	−3.04	0.13	0.052	554	1988–1998
	(618)	(1.13)	(0.24)			
Glamis$^-$	684	−0.89	0.08	0.152	1115	1983–2004
	(100)	(0.15)	(0.12)			
Goldcorp$^-$	339	−0.29	0.18	0.131	1125	1982–2004
	(45)	(0.06)	(0.12)			
Homestake$^-$	3539	−1.84	0.22	0.254	1296	1977–2001
	(368)	(0.31)	(0.09)			
Meridian	3131	−1.07	0.15	0.226	911	1987–2004
	(1307)	(2.47)	(0.14)			
Newmont	23979	−21.22	0.32	0.304	287	1986–1991
	(10694)	(16.52)	(0.17)			
Placer Dome$^-$	35716	−31.32	0.65	0.383	592	1987–1998
	(8165)	(14.63)	(0.12)			
Royal Oak	2097	−2.86	0.17	0.121	395	1991–1999
	(0)	(3.13)	(0.31)			
Vista$^+$	−11	0.46	0.39	0.046	964	1986–2004
	(76)	(0.17)	(0.41)			

$^+$ indicates statistically significant convexity at the 5% level.
$^-$ indicates statistically significant concavity at the 5% level.
GLS with correction for heteroskedasticity caused by S_{it-1} in residual (equation (17)).
White standard errors are in parentheses.

in wealth; in Figure 2, this proportion is $\Delta F/\Delta V$. Even if the firm is fat free when $p_g =$ \$409.15/oz., the aggregate estimate for the 17 firms would imply that at least 5.7% of the potential wealth gain when the price increases to \$501.55 is not realized or at least not passed to shareholders.

□ **Estimation of a piecewise-linear value function.** Estimating a quadratic value function is a natural starting point, because we are interested in the curvature of the relationship, but the quadratic is quite restrictive. As an alternative and a sensitivity test, we also estimated a piecewise-linear relationship between S and p_g, with breaks at the 25th and 75th percentiles of the distribution of p_g.[21] To accommodate tests of slope differences, the regressions are run with a slope term in effect over all prices (Δp_{gt}) and additional slope terms that apply only for prices

[21] For each regression, we used the 25th and 75th percentiles during the time the firm is in the sample, rather than the values shown in Table 1.

TABLE 3 Estimated Change in Slope of S Function When p_g Increases from Its 25th Percentile ($409.15) to Median ($501.55) Price

Company	(1) (%)	(2) (%)	(3) (%)	(4) (%)
Alta	0.0	0.9	5.0	−9.3
	(2.9)	(2.8)	(5.2)	(7.6)
Amax	13.2	10.4		13.2
	(7.7)	(6.8)		(7.7)
Barrick	159.0	−390.9		−18.4
	(117.8)	(−273.0)		(8.3)
Campbell	−5.6	−4.8	−5.2	26.2
	(1.2)	(1.1)	(1.9)	(7.7)
Canyon	−28.7	−24.5		−26.5
	(10.2)	(9.6)		(9.4)
Coeur D'Alene	−6.2	−5.4	−6.0	248.2
	(1.4)	(1.3)	(1.4)	(124.8)
Dakota	41.8	23.5		41.8
	(39.6)	(34.0)		(39.6)
Echo Bay	8.4	4.4	−19.9	50.6
	(4.7)	(4.0)	(6.7)	(6.8)
Getchell	−39.8	−32.5		−39.8
	(14.8)	(12.7)		(14.8)
Glamis	−25.5	−20.9	−19.1	−25.5
	(4.4)	(3.6)	(7.6)	(7.1)
Goldcorp	−12.1	−10.6	−9.2	−10.3
	(2.5)	(2.2)	(3.5)	(9.8)
Homestake	−6.1	−5.4	−5.4	−11.6
	(1.0)	(0.9)	(0.8)	(3.3)
Meridian	−3.7	−4.2		−3.7
	(8.5)	(7.3)		(8.5)
Newmont	−12.8	−8.7		−10.1
	(10.0)	(10.5)		(8.1)
Placer Dome	−12.6	−11.5		−12.6
	(5.9)	(5.0)		(5.9)
Royal Oak	−28.5	−29.9		−28.5
	(31.2)	(28.5)		(31.2)
Vista	24.0	19.4		26.3
	(9.0)	(9.0)		(10.2)

(1) GLS estimation (equation (17)) with White standard errors (from Table 2).
(2) Same as (1) except using 12- to 14-month-out futures price instead of nearest futures price.
(3) Same as (1) except including only observations during 1977–1986.
(4) Same as (1) except including only observations during 1987–2004.

in the lowest (Δp^L_{gt}) and highest (Δp^H_{gt}) quartiles of the gold price distribution. The equation we estimate is[22]

$$\Delta S_{it} - R_{ft} S_{it-1} = \alpha_1 \Delta p_{gt} + \alpha_{2H} \Delta p^H_{gt} + \alpha_{2L} \Delta p^L_{gt} + \alpha_3 S_{it-1}(R_{mt} - R_{ft}) + S_{it-1}\epsilon. \quad (18)$$

The results, shown in Table 4, are consistent with the quadratic estimation. For 6 of the 17 firms, the slope in the lowest quartile of gold prices is estimated to be significantly (at the 5% level) steeper than in the middle range of prices, indicating concavity. In one case, the slope is significantly flatter in the lowest quartile than in the middle range of prices. For 6 of the 17 firms, the slope in the highest quartile of gold prices is estimated to be significantly (at the 5% level) flatter than in the middle range of prices, indicating concavity, whereas it is not significantly

[22] Thus, the estimated coefficient on the highest quartile is $\alpha_1 + \alpha_{2H}$ and the estimated coefficient on the lowest quartile is $\alpha_1 + \alpha_{2L}$.

TABLE 4 Results from Piecewise Linear Regressions

Company	Δp_{gt}^{L}	Δp_{gt}	Δp_{gt}^{H}	CAPM β	R^2	Slope Change* (%)
Alta	102	89	56	0.43	0.163	−24
	(63)	(17)	(31)	(0.14)		
Amax	−1,114	4,347	127	0.40	0.351	+38
	(740)	(373)	(673)	(0.18)		
Barrick⁻	26,681	541	398	0.66	0.160	−97
	(3,226)	(142)	(298)	(0.19)		
Campbell⁻	−15	115	−68	0.63	0.092	−53
	(27)	(19)	(24)	(0.14)		
Canyon⁻	100	143	−136	−0.30	0.204	−97
	(98)	(49)	(50)	(0.19)		
Coeur D'Alene⁻	209	282	−190	0.59	0.287	−81
	(118)	(41)	(52)	(0.14)		
Dakota	−161	208	−95	−0.03	0.091	+144
	(137)	(125)	(129)	(0.39)		
Echo Bay	−2,611	4,142	−2,405	0.04	0.259	+13
	(586)	(545)	(589)	(0.15)		
Getchell⁻	2,355	286	6	0.19	0.085	−89
	(680)	(91)	(125)	(0.22)		
Glamis⁻	500	118	−15	0.03	0.168	−83
	(71)	(20)	(25)	(0.12)		
Goldcorp⁻	2,818	150	−35	0.11	0.218	−96
	(415)	(18)	(24)	(0.11)		
Homestake⁻	9,037	2,293	−1,343	0.26	0.282	−92
	(1,052)	(283)	(343)	(0.08)		
Meridian	773	2,319	511	0.15	0.229	−8
	(951)	(334)	(435)	(0.14)		
Newmont⁻	16,442	8,085	4,382	0.38	0.342	−49
	(2,229)	(1,306)	(2532)	(0.16)		
Placer Dome⁻	6,489	21,477	−8,933	0.65	0.403	−55
	(4,649)	(1,629)	(2,639)	(0.12)		
Royal Oak	628	565	154	0.16	0.125	−40
	(600)	(389)	(462)	(0.31)		
Vista	−12	152	−7	0.33	0.044	+3
	(73)	(56)	(69)	(0.41)		

* Change in slope from bottom to top quartile.
$^{+}$ indicates statistically significant convexity (at 5% level) when comparing slope in top and bottom quartile.
$^{-}$ indicates statistically significant concavity (at 5% level) when comparing slope in top and bottom quartile.
GLS with correction for heteroskedasticity caused by S_{it-1} in residual (equation (18)).
White standard errors are in Parentheses.

steeper in the highest quartile for any firm. An F-test of whether $\alpha_{2H} = \alpha_{2L}$ indicates that the slope is significantly (at the 5% level) smaller in the top quartile than in the bottom quartile for 10 firms (8 of which indicated significant concavity in the quadratic function estimation) and the slope is not significantly different between the quartiles for the remaining 8 firms.

The percentage difference between the estimated slope in the top quartile and the estimated slope in the bottom quartile is presented in the right-hand column of Table 4. The unweighted average of this statistic across the 17 firms is 39% less slope in the top quartile than in the bottom quartile. Thus, again there is strong evidence that for many of these firms, the slope of the S function is greater when gold prices are low than when they are high.

4. Alternative explanations for concavity

■ Having found significant concavity for a number of firms in our sample, we are tempted to infer that these firms do not always maximize profits given the prices they face, and in particular that increases in wealth will be partly dissipated in inefficiency, or at least that is what investors expect. There are, however, a number of potential alternative explanations.

□ **Progressive corporate profits tax.** The progressive corporate profits tax in the United States—broadly, zero tax when the firm has negative earnings and a linear rate of 34%–48% (varying during our sample period) when it has significant positive earnings—might explain some concavity in the $S(p_g)$ function, by making after-tax flow profits a concave function of pre-tax flow profits.[23] To consider an extreme possibility, suppose that at low values of p_g a marginal pre-tax dollar is untaxed, whereas at high levels it is taxed immediately at rate t. Then taxes reduce the slope of the flow-profit function at high gold prices by a factor $1 - t$, while having no effect at low gold prices.

This calculation is misleading, however, because firms can carry forward losses to offset profits. To illustrate starkly, consider another extreme possibility: suppose that (i) positive annual profits are taxed at 48% and negative profits have no tax consequence, (ii) losses always can be carried forward long enough to offset future profits, and (iii) the discount rate is zero. In that case, all firms would pay 48% on their net (over time) profits. Any change in wealth from a change in the price of gold would be taxed at 48% regardless of the level of gold prices; this would lower the slope uniformly but would not affect convexity or concavity.

In fact, the tax code is much more complex. In particular, tax losses can be carried forward only for a limited amount of time, and lose value when carried forward, because of (time) discounting. Still, because a firm can smooth taxable income across years, its marginal tax rate is likely to vary much less than a simple view of the corporate profits tax schedule would suggest. Altshuler and Auerbach (1990) examine the effect of the tax schedule's nonlinearity on the effective marginal tax rate facing corporations.[24] They find that at a time when nominal marginal corporate tax rates varied between zero and 46%, the effective expected marginal tax varied cross-sectionally from 18.9% to 38.6%.

Even if V were linear and if higher tax rates in the Altshuler-Auerbach range were systematically associated with higher gold prices, this could explain at most a proportional decline in after-tax slope from $1 - 0.189 = 0.811$ to $1 - 0.386 = 0.614$: that is, it could at most make the slope at the highest gold price $[0.811 - 0.614]/0.811 = 0.25$, or 25% lower than the slope at the lowest gold price. Comparing the right-hand column of Table 4, which gives the estimated change in slope from the bottom to the top quartile, suggests that even this substantial effect would not explain the strong concavity we find. For all ten of the firms in which the slopes differ statistically significantly between the bottom and top quartile, the slope in the top quartile is more than 25% lower than the slope in the bottom quartile.

But tax convexity seems unlikely to explain even as much as this 25% quasi-bound, for two reasons. First, it would be surprising if a change in gold price were to move a single firm from Altshuler and Auerbach's minimum to their maximum estimated effective tax rate: the reported variation is across all firms in the sample, including agriculture, mining, construction, manufacturing, transportation, trade, and services, so the variation over time for a firm in just one industry is likely to be smaller. Second, as Graham and Rogers (2002) note, the tax code is most apt to make (present-value) after-tax profits concave in pre-tax profits where (flow) profits are teetering near zero. Because we study equity value, not total firm value, the option value of bankruptcy should tend to make equity substantially convex in total firm value in just those cases where taxes would otherwise most influence the shape of the value function.[25]

[23] Tax policies and rates in Canada and Australia are fairly similar to those in the United States, so this discussion applies to nearly all of the tax liabilities of the firms in the sample.

[24] They study all non-financial corporations, not just gold-mining companies. They account for investment tax credits, credits for foreign income, and a number of other complexities of the tax code.

[25] One might perhaps imagine a firm that is not near bankruptcy even though it is teetering near zero flow profits in the medium term. In our sample, three firms filed for Chapter 11 bankruptcy protection. One other failed to meet NYSE minimum capitalization levels and was delisted.

As Graham and Rogers (2002) also discuss, hedging can further mitigate any tendency for the tax code to tax a firm's good years while not subsidizing its bad years. More generally, the firm's maximized after-tax value should be an upper envelope of (perhaps concave) functions of the price of gold, not simply an after-tax version of the upper envelope of straight lines as described in Section 2 above.

Finally, the possible value concavity due to taxes changed significantly after 1986, when a change in the U.S. tax law made it much more difficult to carry losses backward and forward. Columns 3 and 4 of Table 3 show results when the sample period is broken into 1977–1986 and 1987–2004. If the concavity were due primarily to taxes, we would expect to find more concavity in the latter part of the sample. In fact, of the seven firms for which we have enough data to estimate in both periods (we restricted to at least 104 observations in each period), the value function is significantly less concave after 1986 in three cases, significantly more concave in one case, and not significantly different in three cases.

□ **Omission of relevant correlated prices.** Section 2 noted that omitting a price that is nonlinearly related to p_g could create a spurious concave relationship between firm stock market value and the price of gold. For instance, if the 10-year-out futures price of gold were concave in our gold price regressor p_g, our reported results could obtain even if S were convex in the two prices jointly.

Unfortunately, futures prices for gold did not generally exist for delivery more than 2 years in the future during part of our sample period. But if important "implicit" futures prices were concave in p_g, one would expect to see some indication of this in the longest contract for which prices are available throughout our period, which promises delivery 1 year later than does the nearest contract, or 12–14 months in the future.

One could try to test this explanation by including both this more distant futures price and the nearer futures price in the regression. Because these prices are highly correlated, however, doing so increases the standard errors of the estimates so much that the estimated second derivatives become statistically indistinguishable from zero or from the estimates from the regression without the more distant futures price.[26]

An alternative approach, however, produces evidence against this explanation for concavity. If distant futures prices are concave in nearby future prices, it follows that nearby future prices are convex in long future prices.[27] Thus, omitting the nearby future gold price and using only the more distant futures gold price would be omitting a price that is convex in the included price and would *overestimate* the convexity of the stock price function. Column 2 of Table 3 reports the slope changes implied by these estimates and their standard errors using the 12- to 14-month-out futures price. Evidently this substitution makes very little difference in the results, though it might be causing all estimates of second derivatives to be closer to zero. The estimated concavity of the aggregate portfolio implies that an increase in price from the 25th percentile to the sample median price of gold would decrease the slope of the aggregate S function by 9.4%. Although that is smaller than the 11.5% we estimated using the nearest futures price, the difference is much smaller than one standard error of either estimate.

Finally, we can test directly for a nonlinear relationship between the nearby and more distant futures price series. A linear regression of the 12- to 14-month-out futures price on the nearest futures price and the square of the latter yields a *positive* and significant, though quite small, estimated coefficient on the second-order term. Omission of this futures price would thus tend to cause a bias toward finding convexity. Thus, neither of the tests we have carried out indicates

[26] Though gold price movements cannot be distinguished from a random walk, longer-term mean-reverting behavior is very difficult to diagnose, and investor beliefs about mean reversion even more so.

[27] This is true of the actual relationship between these prices, although not necessarily of an estimated statistical relationship.

that the concavity we find is a result of nonlinearity in the relationship between nearby and more distant futures prices of gold.[28]

The other potentially important omitted output price is the price of silver. Most gold producers also mine some silver, because deposits are often colocated. Of the 17 firms in our sample, 6 exhibit a positive and statistically significant first-order effect of silver prices on firm value in a regression with changes in the prices of both gold and silver. We estimated convexities/concavities (in gold price) after adding first- and second-order silver terms to the estimation of equation (17). Although this changed the estimated second-order effects of gold price changes somewhat, it did not affect the finding that the majority are concave in gold price. S is still estimated to be a concave function of the price of gold for 12 of the 17 firms; for 7 of those, the second derivative is statistically significant at the 5% level. Also, for 12 of the 17 firms, our estimates imply that firm value is concave in the price of silver; for 5 of those, the effect is statistically significant at the 5% level.

Omitted input prices could also potentially be important. If the industry faced increasing marginal costs of some input, then potentially this could transfer (rather than dissipate) the rents generated from high gold prices. *Firm*-level increasing marginal cost does *not* have this effect: even if the firm's marginal exploration project is much more expensive than inframarginal projects, it is still true that the firm can continue to do at a higher p_g what it was doing at a lower p_g, so the upper-envelope result still holds. Even if the input-price effect is an industry-level effect, however, we do not believe this is likely to be very important, for three reasons.

First, the industry executives we talked to did not think it plausibly important (although they did suggest that geologists are better paid when gold prices are high).

Second, changes in the price of the firm's own assets would not explain observed concavity. For instance, the price of gold presumably affects the market value of land on which the gold mine is located, but such changes cannot make the S function concave if the firm owns the land, because they do not affect the basic upper-envelope argument that the firm *could* continue to use the same production plan.

Third, and perhaps most important, even if the short-run industry-level supply curve of some inputs (such as geologists) were sharply upward sloping, so that an increase in p_g would make even inframarginal exploration much more costly in the short run, it seems unlikely that the *long-run* supply curve of geologists is so steeply upward sloping as it would need to be to explain our results. Because our dependent variable is the stock market estimate of the present value of profits, effects on current-year or near-term future profits that do not affect further-out profits will have limited effect on our results. This is particularly true in a competitive extractive industry such as gold mining, where cutting output during an input-price spike would not sacrifice long-run total output but only postpone it.

□ **Debt.** Although the theory concerns the overall (asset) value of the firm, our empirical implementation actually tracks the firm's *equity* value. However, as equity is a call option on the underlying assets (equity holders can own the assets by paying the debt), the value of equity is convex in the value of the assets, so if the latter is convex in p_g, so is the former.[29] This shows that the presence of a *given* amount of debt could not falsely generate concave estimated V functions.

A more subtle possibility would be that firms take on different amounts of debt over time, in a way that is correlated with p_g in the sample. In principle, this could create a spuriously concave equity function. However, we believe that in practice the bias would go the other way. When p_g rises, the increase in the firm's asset value, $V'(p_g)$, is divided between debt holders and equity holders. The fraction of the increase that goes to equity holders is the probability $\lambda(p_g)$ that debt

[28] Mean reversion in gold prices would be a special case of an omitted price: the future prices of gold. For an omitted output price to explain concavity, it would have to be concave in the included price. As discussed, we find no evidence to support this.

[29] We thank an anonymous referee for this elegant argument.

holders will be paid off, that is, that the firm will not enter bankruptcy. So we are estimating the slope of $\lambda(p_g)V'(p_g)$ as a function of p_g. Its derivative is $\lambda V''(p_g) + \lambda'(p_g)V'(p_g)$, which is positive (on the assumption of no fat) provided that, as the theory implies, both V' and V'' are positive and provided that $\lambda'(p_g) \geq 0$. In other words, to find spurious concavity of the empirically estimated V for this reason would require, implausibly, that firms take on so much more debt when p_g is high that they are then significantly *more* likely to go bankrupt than when p_g is low.[30]

□ **Changing variance in gold prices.** In Section 2, we noted that the stock market value of a maximizing firm might be concave in p_g, the current price of gold, if the real-option value of gold mining, which increases with the variance of future spot gold prices, were a concave function of p_g. For this option value to reverse the convexity result, we showed that $\mathrm{E}[p_2^2 \mid p_1]$ must be sufficiently concave in p_1.

We addressed this concern by estimating the relationship between the level and the expected future (uncentered) second moment of gold prices. For every observation, we calculated the sample mean of p_g^2 over the next 26 weekly gold price observations (for the next-to-nearest gold futures contract). We then regressed this measure of the *actual* future second moment of the gold price on a constant, p_g, and p_g^2, using 56 observations spaced 6 months apart over the 28-year sample period. We found that actual future second moment of gold prices is estimated to be a significantly (at the 1% level) *convex* function of p_g.

□ **Hedging by gold-mining firms.** Many gold-mining firms trade in the gold futures market in order to hedge the risk associated with gold price movements.[31] Tufano (1996, 1998) finds substantial diversity in hedging among gold-mining firms.[32] He describes two types of financial hedging that are common in the industry: linear strategies, such as selling gold forward, which reduce the firm's overall exposure to changes in gold prices, and nonlinear strategies, which consist largely of buying options, usually put options.

Holding put (or any other) options will convexify S. As for linear strategies, when a gold-mining firm sells gold forward at a fixed price, this of course flattens out the firm's V (and presumably its S) function; in effect, the firm has already sold some gold, and so now owns less of it, reducing the slope of the V (or S) function. But such a linear strategy considered in isolation should not affect concavity or convexity. A *pattern* of linear strategies, however, could in principle concavify S, if firms in our sample effectively owned less gold when gold prices were high than when they were low. Such a correlation would make their V functions flatter at high gold prices than at low gold prices, so that a concave S could be consistent with no fat, that is, $S = V$.[33]

Unfortunately, because Tufano's data covered only a relatively short time span, we cannot infer whether such a pattern happened to occur; and he showed that firms' risk management practices were changing, so it is difficult to infer whether firms pursued *strategies* that would cause such a pattern, that is, hedge more when gold prices are high.[34] Of course, implementation of such a strategy would require knowledge of when prices are higher than they are likely to be in the future, which would be at odds with efficiency in the gold market.

[30] If a firm cannot access external capital, a wealth increase would relax its internal financial constraint. As the firm responds by adding decreasingly profitable projects, this could raise S at a decreasing rate, potentially making S concave. The firms we spoke with, however, did not mention impediments to external capital.

[31] One executive we spoke with said that when banks lend to a gold-mining firm, they often require or prefer this.

[32] Tufano (1998) shows that a mining stock would be *proportionally* less sensitive to the price of gold as the gold price increases, if the firm has no flexibility in its production plan, because (as his equation (2) confirms) the firm value would be linear in the price of gold.

[33] Selling call options would also tend to concavify the firm's value function. There was some mention of this in our interviews with managers, and Tufano mentions it in his work, but it does not appear to be the primary form of hedging among gold-mining firms.

[34] Systematically selling mines when p_g is high and buying when p_g is low could create concavity, but we found no evidence of this behavior. Furthermore, such transfers *among* the firms in our sample could not explain the aggregate concave value of all the firms in our sample.

To address empirically whether hedging patterns might cause concavity, we examined firms that engage in little or no hedging. Peter Tufano provided us a list of firms that engaged in no hedging activities in 1990 or 1992. Because our discussions with industry participants suggested that hedging has become more common over time, we assumed that such a firm did little or no hedging before 1990. Two firms on this list, Coeur D'Alene and Homestake Mining, were also in our dataset for at least 4 years before 1992. For those two firms, we reestimated S using only data from prior to 1992. In each case, the estimated second-derivative terms were negative (concave S) and significant at the 5% level. The estimated proportional declines in the slopes of S when p_g increases from \$409.15 to \$501.55 are 6.2% for Coeur D'Alene and 5.4% for Homestake. These results are consistent with the results reported in Section 4, when we broke the sample at the end of 1986. In the first 10 years of the sample, when hedging was reported to be less common among gold-mining firms, we find strong evidence of concavity for six of the seven firms that are in our sample for that period.

Finally, Tufano and Serbin (1993) report that the average North American gold producer hedged 9.6 months of output at the end of 1991. Our industry sources indicated that even today, firms seldom hedge more than the equivalent of a few years of their production, so most of their expected future production at any time remains unhedged, especially in light of the "replace your output" general rule discussed below in Section 5. Thus, for both theoretical and empirical reasons, hedging practices are unlikely to cause concavity in the $S(p_g)$ function we estimate.

□ **Optimal labor/executive compensation contracts.** For incentive or risk-sharing reasons, it might be optimal to give managers or workers equity or options in the company. Our analysis is unaffected if they hold equity, because the market value of the firm includes all shareholders. But if they hold options, this could concavify the (remaining) value function of the actual shareholders in the firm.

Similarly, if wages and salaries increase more than linearly with p_g as part of an optimal labor contract (explicit or implicit), this could concavify S, because an increasing share of wealth gains from gold price increases would go to workers, rather than shareholders. Indeed, it would do so in a way very like the "fat" mechanism described above, although we might interpret it differently.

It seems very unlikely, however, that executive compensation tied to earnings could account for a significant fraction of the concavity we find in a substantial number of our sample firms. Gold-mining companies pay a small fraction of firm value as executive compensation.[35] This is not surprising, because a comparatively large share of firm value is represented by tangible, transferable assets. That is, much of the firm value is due to its holdings of land or rights to mine, not value creation by the firm's operations. Furthermore, firm value *changes* are largely due to events (in this case, gold price shocks) that are exogenous to the firm. Incentive/compensation theory suggests that optimal compensation plans should not award managers a significant share of firm value changes that result from plainly observable exogenous events.[36]

To examine this possible explanation, we obtained executive compensation data for the five companies in our sample that are also in Compustat's ExecuComp database.[37] We constructed the total salary and bonus compensation (the "TCC" variable in ExecuComp) and the total

[35] For the five firms for which we obtained recent compensation data, salary plus bonus of the top-five (by salary) employees averaged 0.77% of sales and 0.25% of the equity value of the firm. Total compensation including value of new options issued averaged 2.3% of sales and 0.5% of firm value.

[36] Milgrom and Roberts (1992) call this the "informativeness principle." See, however, Bertrand and Mullainathan (2001), for evidence that managers do get rewarded for luck.

[37] The companies and the years they are in ExecuComp are Getchell (1995–1997), Echo Bay (1994–2001), Barrick (1994–2003), Placer Dome (1994–2004), and Glamis (1994–2003). ExecuComp begins in 1992, but the first 2 years of data are considered less reliable.

compensation including option grants (the "TDC1" variable in ExecuComp[38]) for the top five executives in the firm in each year.[39] This yielded 42 company-year observations.

A log-log regression of the more expansive compensation variable (TDC1) on the firm fixed effects and the equity value of the firm exhibits the expected positive relationship, an estimated elasticity of 0.49 with a White heteroskedastic-consistent standard error of 0.13. The same regression using the price of gold instead of firm value, however, yields a coefficient and standard error of −0.21 (0.45). As compensation theory would suggest, it appears that gold-mining executives are not rewarded for gold price changes. When both firm value and gold price are included in the regression, the effect of gold price is significantly negative and firm value is significantly positive, suggesting that holding the value of the firm constant, compensation declines when the price of gold increases. These results are consistent with the idea that managers are compensated for the value of firm changes *not* driven by gold prices, but not for firm value shocks due to gold price changes.[40] They suggest that it is very unlikely that executive compensation explains a concave relationship between gold prices and the equity value of the firm.

Mining labor costs are a much larger share of firm operating costs than executive compensation. In some industries, labor rent sharing has been suggested as a substantial effect when firms get wealthy. We discuss this in Section 5.

□ **Environmental liabilities.** Gold-mining firms, which are viewed as causing extensive environmental damage, might be required to pay disproportionately more for cleanup if they are relatively rich. Although environmental liabilities are nontrivial—one source put them at about 15% of "hard" costs—this is unlikely to explain the concavity we observe. According to our industry and government sources, most environmental legislation bearing on mining companies applies to all mining, not to specific sectors such as gold mining. Industry participants did not see cleanup costs or liabilities as being very much subject to discretion or variation. Relatedly, when we mentioned this hypothesis to government regulators, they commented that they were aware of no examples. Finally, for environmental liabilities short of bankruptcy to explain concavity, they would have to not only increase with but also be convex in the price of gold, that is, the proportion of marginal wealth that would be allocated to additional environmental liability would have to increase with the price of gold. Finally, one could also ask about "asbestos-style" liabilities, which with some probability will bankrupt the firm. For such risks to cause concavity, however, the probability of a bankrupting liability would have to be significantly increasing in p_g.

□ **Royalty payments.** Sometimes governments (or owners of auriferous properties who delegate the mining) demand royalties for gold extraction. A linear royalty schedule (whether on units, revenues, or profits), like a linear tax schedule, would not affect the predictions of convexity. Royalty rates that increase with the price of gold (or with the total revenues attributable to a mine, for instance), however, could potentially make a fat-free value function concave. Our discussions with industry and government contacts suggested that royalties are most often linear, though. Some royalties kick in above a certain point, and others are capped; thus, some would contribute to concavity and others tend toward convexity. Some royalties are based on accounting net profits, but one well-informed commentator suggested that it is viewed as unwise to take a percentage of the net, because doing so stimulates cost accountants' creativity in undesirable ways, somewhat as it is said to do in the case of Hollywood movies.

[38] Defined by ExecuComp as "Salary, Bonus, Other Annual, Total Value of Restricted Stock Granted, Total Value of Stock Options Granted (Using Black-Scholes), Long-Term Incentive Payouts, and All Other Total."

[39] Glamis lists only four executives for some of the years, so we took the sum of compensation for the top four listed by Glamis in each year.

[40] Using TCC, which includes only salary and bonus, as the dependent variable, the effect of firm value alone is insignificant positive and gold price alone is insignificant negative. With both regressors included, firm value again is positive and significant and gold price again is negative and significant.

We also note that the hypothetical examples this source used in discussing the matter with us had royalty rates of 1% or a few percent, except for one that was 10% of accounting profits. Other industry sources also tended to come up with examples involving a few percentage points. This itself suggests that royalties as a whole are unlikely to be enough to drive our results.[41]

□ **Negative correlation between gold price and gold reserves.** The theory above, together with a Hotelling version of the random-walk theory of gold prices, implies that $V'(p_g) = e(p_g)$, where $e(p_g)$ is the firm's economic reserves of gold.[42] The optimizing theory that we test assumes that e is weakly increasing in p_g. Although higher gold prices clearly make more gold economic to extract, this causal effect could be obscured if other factors induce a negative correlation between economic reserves and gold prices.

Such negative correlation could be just a fluke during our sample period: these companies might have happened to expand their reserves, either through new discoveries or purchase of other companies, at times when gold prices were declining. Less coincidentally, an industrywide improvement in exploration or extraction technology could increase the economic reserves of a typical firm, and also lower gold prices through an increase in (expected) market supply, though it seems unlikely that incremental changes in extraction technology will have a discernible effect on price. Or, when prices rise, mining firms might increase extraction of gold even more than they increase discoveries of new economic reserves, which would create a negative correlation of reserves (on which the company has claim) and prices.

In fact, the real price of gold trended downward from 1980 to 1999, which includes most of our sample period. The question then is whether economic reserves were moving inversely to the gold price trend. Unfortunately, we were unable to find consistent data on economic reserves by firm, but we combined a number of sources to get annual reserve figures for 38 firm-years covering six firms. Using these figures, we estimated (17) adding an additional term: $\alpha_4 reserves_{it}$ Δp_{gt}. This term allows the slope of the S function to vary with the reserves firm i holds at time t.[43]

When we estimate this modified equation by firm, the second-order term is not significant for any of the firms. This seems due in part to the small number of observations and in part to the fact that reserves for most of the firms we observe move very little during the years we observe. As a result, the new variable is highly collinear with Δp_{gt}. A test for pooling the observations across firms, however, does not reject pooling when the reserves interaction variable is included. In a pooled regression, the second-order term is negative and significant at the 6% level.

Finally, we also examined reported U.S. and world economic gold reserves, though we were cautioned that these figures are not very reliable. Neither exhibited the negative correlation with prices over our sample period that might suggest the concavity is being driven by such a statistical artifact.

5. Where's the fat?

■ If indeed fat explains a substantial fraction of the concavity of S that we often find, what is this fat? And how do firms vary in the extent to which they are subject to such fat: for instance, might it be related to size, or to the absence of large shareholders?

Unfortunately, our ability to study such questions is very limited, both because we have only 17 firms and, more fundamentally, because our methodology yields bounds, not estimates,

[41] Finally, Centurion is a company that leases to mining companies properties with known gold deposits and collects royalties. If it imposed more-than-linear royalties, one would expect its $S(\cdot)$ function to be convex. In fact, our estimate for Centurion was concave (though not statistically significant).

[42] This is the first derivative of the value function, which is unaffected *locally* by the option-value effect. Still, the value of e would have to be adjusted for "anticipated finds" if some gold exploration locations have strictly positive expected net present value on further exploration.

[43] More precisely, we used the firm's reported economic reserves for the year in which the observation occurs. Conversations with industry participants, however, did not encourage us to put a lot of faith in these firm-level reports.

for fat. Although we find statistically and economically significant variations in the curvature of different firms' S functions, we would expect different firms to have markedly differently curved V functions, and, as we do not observe V, we cannot confidently infer anything about variations in the fat functions.

Nonetheless, a natural response to our findings is to ask wherein this fat consists, and what determines how much of it there is. We asked these questions in our interviews with industry executives. Below, we discuss two places we have looked for fat and for factors affecting the extent of fat.

□ **Exploration costs.** All the managers we spoke with seemed to believe that—either as an obviously sound business policy or because of pressure from stock market analysts—a gold-mining firm should "replace" its extraction, whether by exploration for new reserves or by acquisition of existing mines (or of mine-owning companies). Several suggested that when gold prices are high, firms found themselves "having to" and/or "able to" undertake quite unpromising exploration projects.

Because it is much harder to verify whether an exploration decision is value increasing than whether a mine is being well managed, exploration seems a likely locus for potentially value-reducing expenditures. In related work in the oil industry (Borenstein and Farrell, 1996), oil industry commentators told us that the industry dissipated much of the value increase during the early 1980s by "excessive" (at least *ex post*) exploration. Clearly, a price increase should induce some increase in exploration, but these observers suggested that the oil industry's response was excessive.[44]

In gold mining, a general rule that firms must replace extraction would suggest one simple principal-agent theory for value dissipation after gold price increases. Suppose that mines are run as profit centers, or more broadly that mine managers have incentives to increase output when p_g rises, in a way that takes account of increased *extraction* costs but does not take account of the marginal cost of finding more gold. Then their output-increasing decisions, although optimal if the firm optimized overall, could actually reduce the firm's value if the firm forced itself to follow the general rule that it must replace all extraction.

Another possible theory, attributing the anomaly to the financial markets rather than to a principal-agent problem inside the firm, would be that some firms resist this general rule and are penalized by stock market analysts. Several executives told us that they believe analysts behave in this way.

□ **Non-optimal labor compensation.** If the firm optimally increases labor and compensation as p_g rises, this would not lead to concavity of the value function. To the extent that this goes beyond an optimal *ex ante* contract, however, and becomes an inefficient *ex post* holdup or asset stripping, one might call it fat and it could make the value function concave. Though it can be hard to distinguish efficient from inefficient variations in labor compensation, it seems unlikely to be efficient to reward miners for changes in firm value driven by exogenous changes in the price of gold. Although such labor rent sharing has been documented in some industries,[45] our discussions with industry participants suggested that it is not likely to be much of an issue in gold mining. None reported that wage rates moved noticeably with the price of gold.

6. Conclusion

■ Once one recognizes that firms could be inefficient, one might suspect that they get fatter as their wealth grows. In the gold-mining industry, we found empirically that an increase in

[44] Jensen (1986) presents evidence of value-reducing exploration in the oil industry.

[45] See, for instance, Rose (1987).

gold prices increases many firms' stock market values by more when the price of gold is low than when it is already high. This empirical concavity contradicts the basic theoretical prediction of convexity driven by the upper-envelope, or real-option, effect for an optimizing firm.

The concavity result is particularly striking in that real options are important in gold mining. As Moel and Tufano (2002) document, firms open and close mines in response to changes in the price of gold. Standard theory suggests that such flexibility in production plans should make a firm's value strongly convex as a function of the price of gold. We find that for more than half of the gold-mining firms we study, it does not.

We posit that much of the observed concavity reflects investors' beliefs that firms will dissipate a share of wealth gains and that this share will be larger when the firm is wealthy. We discuss and reject a number of alternative explanations for the concavity we find.

Gold mining is a price-taking industry (and our study relied on this); thus, variation in market power is not the source of the wealth variations we study. In other industries, however, market power can be a source of firm wealth. Our results are consistent with the popular view that the resulting fat may dissipate what would otherwise be monopoly profits, increasing (perhaps dramatically) the deadweight loss of monopoly. If so, competition may improve productive as well as allocative efficiency.

References

ALTSHULER, R. AND AUERBACH, A.J. "The Significance of Tax Law Asymmetries: An Empirical Investigation." *Quarterly Journal of Economics,* Vol. 105 (1990), pp. 61–86.

BERTRAND, M. AND MULLAINATHAN, S. "Are CEOs Rewarded for Luck? The Ones without Principals Are." *Quarterly Journal of Economics,* Vol. 116 (2001), pp. 901–929.

BORENSTEIN, S. AND FARRELL, J. "Is There Fat in Oil? Diagnosing Profit Dissipation from the Stock Market Value of Oil Companies." Manuscript, University of California, Berkeley, 1996. Available at faculty.haas.berkeley.edu/borenste/oilfat.pdf.

GORMAN, W.M. "Measuring the Quantities of Fixed Factors." In J.N. Wolfe, ed., *Value, Capital, and Growth: Papers in Honour of Sir John Hicks*. Edinburgh: Edinburgh University Press (Chicago: Aldine Publishing), 1968.

GRAHAM, J.R. AND ROGERS, D.A. "Do Firms Hedge in Response to Tax Incentives?" *Journal of Finance,* Vol. 57 (2002), pp. 815–839.

JENSEN, M.C. "Agency Costs of Free Cash Flow, Corporate Finance, and Takeovers." *American Economic Review Papers and Proceedings,* Vol. 76 (1986), pp. 323–329.

KREPS, D.M. *A Course in Microeconomic Theory*. Princeton, NJ: Princeton University Press, 1990.

LEIBENSTEIN, H. "Allocative Efficiency versus X-Efficiency." *American Economic Review,* Vol. 56 (1966), pp. 392–415.

MCFADDEN, D. "Cost, Revenue and Profit Functions: A Cursory Review." Working Paper no. 86, Institute of Business and Economic Research, University of California, Berkeley, 1966.

MILGROM, P. AND ROBERTS, J. *Economics, Organization and Management*. Englewood Cliffs, NJ: Prentice-Hall, 1992.

MOEL, A. AND TUFANO, P. "When Are Real Options Exercised? An Empirical Study of Mine Closings." *Review of Financial Studies,* Vol. 15 (2002), pp. 35–64.

PINDYCK, R.S. "The Present Value Model of Rational Commodity Pricing." *Economic Journal,* Vol. 103 (1993), pp. 511–530.

ROSE, N.L. "Labor Rent Sharing and Regulation: Evidence from the Trucking Industry." *Journal of Political Economy,* Vol. 95 (1987), pp. 1146–1178.

SCHMITZ, J.A., JR. "What Determines Productivity? Lessons from the Dramatic Recovery of the U.S. and Canadian Iron Ore Industries Following Their Early 1980s Crisis." *Journal of Political Economy,* Vol. 113 (2005), pp. 582–625.

SELVANATHAN, E.A. "A Note on the Accuracy of Business Economists' Gold Price Forecasts." *Australian Journal of Management,* Vol. 16 (1991), pp. 91–95.

TUFANO, P. "Who Manages Risk? An Empirical Examination of Risk Management Practices in the Gold Mining Industry." *Journal of Finance,* Vol. 51 (1996), pp. 1097–1137.

———. "The Determinants of Stock Price Exposure: Financial Engineering and the Gold Mining Industry." *Journal of Finance,* Vol. 53 (1998), pp. 1015–1052.

——— AND SERBIN, J. "American Barrick Resources Corporation: Managing Gold Price Risk," Harvard Business School Case no. 9-293-128, 1993.