

Growth versus Margins: Destabilizing Consequences of Giving the Stock Market What It Wants

PHILIPPE AGHION and JEREMY C. STEIN*

ABSTRACT

We develop a model in which a firm can devote effort either to increasing sales growth, or to improving per-unit profit margins. If the firm's manager cares about the current stock price, she will favor the growth strategy when the market pays more attention to growth numbers. Conversely, it can be rational for the market to weight growth measures more heavily when it is known that the firm is following a growth strategy. This two-way feedback between firms' strategies and the market's pricing rule can lead to excess volatility in real variables, even absent any external shocks.

IN HIS COMPANY'S ANNUAL REPORT FOR 2001, Amazon.com founder and CEO Jeffrey Bezos begins his letter to shareholders as follows:

In July of last year, Amazon.com reached an important way station. After 4 years of single-minded focus on growth, and then just under 2 years spent almost exclusively on lowering costs, we reached a point where we could afford to balance growth and cost improvement, dedicating resources and staffed projects to both.

As this statement suggests, many firms face a fundamental strategic tradeoff: They can focus their efforts on increasing sales growth or on improving profit margins (e.g., by lowering unit costs), but given limits on managerial time and other resources, doing more on one dimension necessarily implies doing less on the other. In other words, the choice of whether to emphasize growth or margins is essentially a multitasking problem of the sort described by, for example, Holmstrom and Milgrom (1991).

In this paper, we study the implications of this multitasking problem for the dynamics of firm behavior. We are particularly interested in the interplay between firms' strategies and the way that the stock market goes about evaluating them. Anecdotal evidence suggests that, just as a firm can change its strategic orientation over time, so too can investors shift the emphasis that they place on

*Aghion is from Harvard University and the National Bureau of Economic Research. Stein is from Harvard University and the National Bureau of Economic Research. We are grateful to the National Science Foundation for financial support, and to Richard Holden and Naomi Hausman for research assistance. Thanks also to an anonymous referee, Malcolm Baker, Patrick Bolton, Arvind Krishnamurthy, Tuomo Vuolteenaho, and seminar participants at the NBER, NYU, and Boston University for helpful comments.

various performance measures. In the specific case of Amazon.com, Hong and Stein (2003) provide a detailed reading of equity analysts' reports and document that, during the period that Bezos calls the growth phase (roughly through the end of 1999), analysts were almost uniformly focused on growth-related metrics in valuing Amazon stock, to the virtual exclusion of profitability or cost-related metrics.¹ Conversely, during the cost-cutting phase that followed, analysts began to pay much more attention to per-unit measures of costs and profits.²

One interpretation of the Amazon story is that managers actively cater to the stock market's preferences in choosing their strategies. That is, when the market is more interested in growth than profit margins, managers like Bezos give the market what it wants by focusing their efforts on delivering increased revenue growth, and when the market becomes more interested in margins, managers take the cue and adapt their strategies accordingly. Certainly, if one thinks of the market's valuation model as exogenously given, it would seem to make sense for a manager interested in maximizing her firm's current stock price to cater in this fashion.

Alternatively, one might read the Amazon story as saying that firms' strategies drive the market's valuation model, not the other way around. And again, it would seem perfectly rational for investors to pay more attention to a firm's growth numbers if they know that management is devoting most of its effort to generating growth.³

We take the view that the causality runs in both directions. In our model, managers interested in their stock prices do indeed cater to the market, but at the same time, the market's valuation model rationally takes into consideration what it perceives to be a firm's current business strategy. This two-way feedback can produce a variety of interesting effects. Perhaps most notably, it can lead to autonomous fluctuations in output growth, even absent any external source of shocks.

The first step in understanding how these fluctuations arise is to note that in a static setting, the two-way feedback can generate multiple equilibria when parameter values are in an "intermediate" range such that neither a growth-oriented nor a margins-oriented strategy represents too extreme a deviation from first-best efficiency. For example, if the market conjectures that the firm is pursuing a growth strategy, its valuation will tend to put more weight on realized growth, which will in turn encourage the manager to stick with the growth strategy so long as this is not too inefficient. On the other hand, if

¹ For example, in a February 1999 report lauding Amazon's growth prospects, Ehrens (1999) of Bear Stearns writes that: "In traditional off-line retail, the gross margin is a very important metric to watch . . . This is not the case in the on-line world."

² In a March 2001 report on Amazon, Becker (2001) of Lehman Brothers argues: "It is now clear that higher customer churn rates, weak shipping margins and equally high marketing spend will offset many of the company's virtues . . . we recommend investors stay on the sidelines . . . Clearly, the company will need to increase gross margins . . ."

³ As Dewatripont, Jewitt, and Tirole (1999) argue, and as we develop in more detail below, this is a natural feature of the sort of learning model introduced by Holmstrom (1999). Intuitively, investors can learn more about a manager's general ability by looking at those performance measures that she is most actively trying to maximize.

the market conjectures that the firm is pursuing a cost-cutting strategy, its valuation will tend to emphasize margins, which will reward the manager for staying with the cost-cutting strategy.

Now consider the dynamics. Imagine a firm that starts out endowed with such attractive growth prospects that, irrespective of the market's conjecture, the optimal choice is for it to pursue a growth strategy. This puts us initially in the growth equilibrium, in which case the market price is especially sensitive to performance on the growth dimension. Over time, as the firm penetrates the market more fully, it will begin to find the growth strategy less attractive. More specifically, it will eventually reach a point where, if it were only interested in choosing the first-best action, it would start shifting resources away from the growth strategy and toward the margins strategy. But here is where things get difficult. For as long as the market continues to value it as a growth firm, any change in strategy will lead it to disappoint the market on the growth dimension, thereby damaging its stock price. Thus, if it cares about pleasing the market, the firm will be trapped into continuing with the growth strategy longer than is optimal, instead of attending to cost-cutting, as it should.

Eventually, the market will become so saturated, and the growth strategy so inefficient, that the firm will have no choice but to switch to a margins strategy. However, relative to the first-best, this shift will not only be too late, it will be too abrupt. Much like Bezos's description of Amazon post-1999, the firm will go to the other extreme of focusing exclusively on cost-cutting, as opposed to taking a balanced approach of devoting some resources to each of the two strategies. This lack of balance in turn sets the stage for another round of fluctuations. Once entrenched in the margins equilibrium, with the market now expecting strong performance on the margins dimension, the firm will pay insufficient attention to growth opportunities for too long, and will eventually get to a point where it is forced to go back to the growth strategy, at which point the whole process begins again.

The bottom line is that unlike in the first-best scenario, where the firm responds to a smooth decline in its growth prospects by gradually and monotonically shifting resources away from a growth strategy and towards a margins strategy, adjustment in a world in which managers cater to the market involves a series of sharp oscillations. The firm may go from an all-growth strategy, to an all-margins strategy, then back again to an all-growth strategy, before eventually settling down. Central to this result is the premise that managers care to some degree about maximizing their current stock price, as opposed to just the present value of future cash flows. Indeed, the more intensely a manager cares about the current stock price, the more dramatic are the associated fluctuations.

The idea that the stock market can lead to excessive volatility in investment and output is a venerable one, going back at least to Keynes (1936). More recent treatments of this idea have tended to emphasize the role of irrational noise traders, who create exogenous deviations of stock prices from fundamentals. The question typically asked is to what extent these mispricings influence

corporate investment, either at the micro or macro level.⁴ Our approach in this paper is quite different. In most of our analysis, we model the market as fully rational, albeit imperfectly informed about managerial ability.⁵ Moreover, beyond this uncertainty about ability, we introduce no further exogenous sources of variation, that is, there is no analog in our setting to noise trader-induced sentiment shocks. Thus, the real-side volatility in our model is driven by intrinsic forces, and changes over time in the market's approach to valuation are completely endogenized. As we argue in detail below, this leads to a novel set of empirical implications.

The remainder of the paper is organized as follows. In Section I, we introduce a static version of our model. This static version leaves a lot out; in particular, it takes the demand for a firm's products to be exogenous. However, it helps to develop the basic intuition for how multiple equilibria can arise in our setting. In Section II, we study dynamics. In so doing, we endogenize the demand curve facing the firm at any point in time via a market saturation mechanism of the sort alluded to above: When the firm plays the growth strategy for several periods in a row, this reduces the stock of unsatiated customers and hence shrinks demand going forward. In Section III, we flesh out the model's most distinctive empirical implications. Section IV discusses related work, and Section V concludes.

I. Static Model

A. Technology and Information Structure

To build intuition and lay the groundwork for our dynamic model, we begin with a very simple static setup. We consider a firm that produces for two periods. In the first period, the firm's manager must decide how to allocate her effort between two business strategies: (i) a "growth" strategy of expanding production and sales; and (ii) a "margins" strategy of maintaining sales while improving profit margins (say by reducing unit costs). Following Holmstrom (1999), we assume that the manager's ability is initially unknown to both the manager herself and the market. However, her ability is inferred by the market at the end of the first period, based on the firm's performance at this time. Managerial ability is captured by a parameter a that is normally distributed with mean A and variance v^a .

The firm's first-period profit is the sum of two components, a sales-volume component s_1 and a margins component m_1 .⁶ Thus,

⁴ See, for example, Bosworth (1975), Fischer and Merton (1984), Morck, Shleifer, and Vishny (1990), Blanchard, Rhee, and Summers (1993), Stein (1996), Chirinko and Schaller (2001), Baker, Stein, and Wurgler (2003), Polk and Sapienza (2006), and Gilchrist, Himmelberg, and Huberman (2005).

⁵ In contrast to our theory, most papers that assume a rational, information-aggregating stock market have tended to come to positive conclusions regarding its consequences for economic efficiency. See, for example, Holmstrom and Tirole (1993).

⁶ Naturally, it would be more realistic to write profits as the product of sales and a profit margin, rather than as the sum. While such a multiplicative formulation is a bit more cumbersome than our additive simplification, it leads to all the same basic results.

$$\pi_1 = s_1 + m_1, \quad (1)$$

where s_1 and m_1 depend on both managerial ability as well as the manager's allocation of effort.

The manager is endowed with one unit of effort, a fraction e of which is devoted to the growth strategy of increasing s_1 and the remaining $(1 - e)$ of which is devoted to improving margins m_1 . The manager chooses effort so as to maximize the expectation of her utility U , which we take to be equal to a linear combination of current profits, the firm's current stock price, and discounted future profits,

$$U = \pi_1 + \alpha P_1 + (1 - \alpha)\delta\pi_2, \quad (2)$$

where P_1 is the stock price at the end of the first period (after π_1 has been paid out as a dividend), π_2 is the second-period profit, and δ is a discount factor.

This sort of managerial utility function, which appears in a number of previous papers (e.g., Miller and Rock (1985) and Stein (1989)), corresponds to the assumption that the manager comes into the first period owning shares of the firm's stock, and plans to sell a fraction α of these shares at the end of the first period while keeping the remaining $(1 - \alpha)$ through the end of the second period. This policy yields the manager a first-period payment proportional to $(\pi_1 + \alpha P_1)$ —the first term reflecting the dividend on her shares, and the second reflecting the ex-dividend proceeds on those shares she sells—and a second-period payment proportional to $(1 - \alpha)\pi_2$ per share, which gets discounted by δ . It turns out that we can disregard the $(1 - \alpha)\delta\pi_2$ term in the utility function because as we will see momentarily, the manager's choices of actions have no consequence for second-period profits. Thus, for the purposes of analyzing her behavior, the manager's utility function boils down to

$$U = \pi_1 + \alpha P_1. \quad (3)$$

Our specification of managerial utility raises an important question. Since the whole point of this paper is to highlight the distortions that arise when managers care about current stock prices, that is, when $\alpha > 0$, one has to ask why the firm would ever design an incentive scheme with $\alpha > 0$ in the first place. In other words, even if the firm has to provide the manager with shares of stock for incentive purposes, why not make this stock restricted, so that it cannot be sold until the end of the second period?

We provide a formal treatment of this question in Appendix A, extending the model so as to endogenously deliver a compensation scheme that maps exactly into our assumed managerial utility function. We do so by making two further assumptions. First, there is an additional moral hazard problem, in that the manager views it as costly to provide basic maintenance of the firm's assets. This moral hazard problem creates the need for managerial stock ownership.

Second, the manager sometimes experiences a liquidity shock similar to that in Diamond and Dybvig (1983), and in such states of the world can only derive utility from payouts received in the first period. This implies that if the firm is to derive the full incentive benefits of stock ownership, it cannot make the

stock entirely restricted, rather, it has to allow the manager to sell the stock early when she experiences the liquidity shock.

In spite of the distortions that this entails, making the manager's stock unrestricted in this fashion can still be on net worthwhile for the firm, to the extent that the moral hazard problem associated with asset maintenance is sufficiently consequential.⁷ It is also worth emphasizing that unrestricted equity compensation of the sort that our model envisions is the norm in the real world. For example, Malmendier and Tate (2005) document that the typical executive stock option grant is composed of options with a 10-year maturity that are fully vested and can be freely exercised—and, indeed, often are—after 4 or 5 years.

We assume the following technologies for generating first-period sales and margins, respectively:

$$s_1 = \alpha e q_1 + \varepsilon^s, \quad (4)$$

and

$$m_1 = \alpha(1 - e) + \varepsilon^m, \quad (5)$$

where: (i) q_1 denotes the size of the market, which for the time being we take to be exogenously fixed (an assumption we relax in the next section when we analyze the dynamics of the model); (ii) ε^s is a random variable that captures sales shocks and is normally distributed with mean zero and variance v^s ; and (iii) ε^m is a random variable that captures margins shocks, is normally distributed with mean zero and variance v^m , and is independent of ε^s .

The technologies for sales and margins embody two key assumptions. First, ability and effort are complements when it comes to improving either aspect of performance. As emphasized by Dewatripont et al. (1999), this implies that one learns more about an agent's ability by looking at her performance on the dimension on which she is working the hardest.⁸ Second, there is an asymmetry across the two tasks: Both managerial effort and ability have a higher marginal product in terms of increasing sales when the size of the market q_1 is larger, but their impact on margins is independent of market size. This means that, insofar as the manager is interested in maximizing current profits, she will be more inclined to devote her efforts to the growth strategy when q_1 increases.

⁷ There are various other ways to endogenize a managerial incentive contract of the sort we assume. Laffont and Tirole (1988) develop a model in which an analog of our parameter α corresponds to the (optimally-chosen) probability that the incumbent manager will be taken out of her position in the stock by a hostile takeover. In Bolton, Scheinkman, and Xiong (2006), firm founders choose to give managers short-horizon stock-based incentives because the founders themselves plan on selling out soon, and hence want to induce managers to care about the current stock price, even if this leads to distorted investment.

⁸ To take a canonical example: A high school student's inability to speak French can only say something negative about her broader capabilities if she has been enrolled in French classes; if she has been studying Spanish instead, no inference can be drawn.

In the second period, we assume that the firm's profits are given by

$$\pi_2 = a + \rho(\varepsilon^s + \varepsilon^m), \quad (6)$$

where $\rho \in (0, 1)$ measures the persistence of the sales and margin shocks over time. Thus, not only do the sales and margins shocks partially carry over from one period to the next, but the firm is also better off in the second period, all else equal, with a higher-ability manager. Note, however, that managerial *actions* have no consequence for second-period profits, as mentioned above.

The timing within the first period is as follows. First, the manager chooses an allocation of effort; this effort level cannot be directly observed by the market.⁹ Second, sales and margins are realized and publicly observed. Third, the market updates its forecast of second-period profits, based on the observed values of sales and margins, and its prior conjecture as to how the manager allocates effort. This in turn determines the first-period stock price P_1 . The stock price is simply the discounted expectation of second-period profits, which implies that

$$P_1 = \delta E(\pi_2 \mid s_1, m_1, e^*), \quad (7)$$

where e^* is the allocation of managerial effort that the market conjectures, and δ is the market's discount factor between the first and second periods.

B. First-Best Action

The natural benchmark is the case in which the manager has no incentive to distort her effort allocation in order to impress the market. This case arises when either the manager's ability a is public information *ex ante* (i.e., when $v^a = 0$), or when the manager does not care about the stock price (i.e., when $\alpha = 0$). If either of these conditions holds, the manager will simply choose e so as to maximize expected first-period profits:

$$\max_e \{E(\pi_1) = aeq_1 + a(1 - e)\}. \quad (8)$$

It then immediately follows that the static first-best action is to set $e = 1$ (i.e., to pursue a growth strategy) if $q_1 > 1$ and to set $e = 0$ (i.e., to pursue a margins strategy) if $q_1 < 1$. Again, the intuition for this result is that the marginal return on effort devoted to the growth strategy is greater when the size of the market, q_1 , is larger.

⁹ A key assumption is that when the manager entertains a deviation in strategy, for example, from growth to cost-cutting, she presumes that the market will not immediately recognize this deviation. This should not be taken as saying that investors in the real world do not often have a good understanding of firms' strategies. Two points are relevant in this regard. First, in equilibrium the market always makes a rational conjecture at time t of the firm's time- t strategy, even if it doesn't observe this choice contemporaneously. Second, in the dynamic version of the model that we develop below, the market can always observe the firm's *past* strategy choices, that is, it can directly observe the time- t strategy at time $t + 1$. Since in our dynamic equilibrium these strategies tend to be highly persistent, this will resemble a situation in which the market "knows" what the firm is going to do at time $t + 1$ because it has observed the strategy it is following at time t .

C. The Interaction between the Firm and the Market

To solve the model in the more general case in which the manager does try to influence the market's perceptions, we proceed as follows. We begin by endowing the market with a particular conjecture about managerial actions (e.g., the market might conjecture that the manager is pursuing a growth strategy). Given this conjecture, we can solve for market prices as a function of the firm's observed performance on the growth and margins dimensions. We then take the manager's perspective and ask what strategy she will prefer to follow in light of the market pricing rule. In a rational expectations equilibrium, the manager's optimal strategy coincides with the market's conjecture. In other words, an equilibrium is a pair (e, e^*) such that: (i) anticipating the market conjecture e^* , the manager chooses an effort allocation e ; and (ii) e confirms the market conjecture, namely: $e = e^*$.

C.1. Stock Prices

Suppose first that the market conjectures that the firm is pursuing a growth strategy ($e^* = 1$). Then, based on the observed realizations of s_1 and m_1 , Bayesian updating leads to the following pricing rule for the firm's stock:

$$\begin{aligned} P_1^g &= \delta E(\pi_2 | s_1, m_1, e^* = 1) \\ &= \delta(A + \beta^{sg}(s_1 - Aq_1) + \beta^{mg}m_1), \end{aligned} \quad (9)$$

where, using the fact that the two variables s_1 and m_1 are uncorrelated, we have that

$$\beta^{sg} = \text{cov}(\pi_2, s_1 | e^* = 1) / \text{var}(s_1 | e^* = 1) = (q_1 v^a + \rho v^s) / (q_1^2 v^a + v^s) \quad (10)$$

and

$$\beta^{mg} = \text{cov}(\pi_2, m_1 | e^* = 1) / \text{var}(m_1 | e^* = 1) = \rho. \quad (11)$$

Similarly, suppose that the market conjectures that the firm is pursuing a margins strategy ($e^* = 0$). Then, based on the observed realizations of s_1 and m_1 , Bayesian updating leads to the following pricing rule:

$$\begin{aligned} P_1^m &= \delta E(\pi_2 | s_1, m_1, e^* = 0) \\ &= \delta(A + \beta^{sm}s_1 + \beta^{mm}(m_1 - A)), \end{aligned} \quad (12)$$

where

$$\beta^{sm} = \text{cov}(\pi_2, s_1 | e^* = 0) / \text{var}(s_1 | e^* = 0) = \rho \quad (13)$$

and

$$\beta^{mm} = \text{cov}(\pi_2, m_1 | e^* = 0) / \text{var}(m_1 | e^* = 0) = (v^a + \rho v^m) / (v^a + v^m). \quad (14)$$

A couple of features of these pricing rules should be emphasized. First, we have that $\beta^{mm} > \beta^{sm}$. In other words, when the market thinks the manager is

playing the margins strategy, it puts more weight on margins than on sales. This makes intuitive sense: When the manager is actively trying to improve margins, the realization of m_1 is more informative about her ability, and hence about future profits, than is the realization of s_1 . Second, if $\rho q_1 < 1$ (which will always be satisfied in a neighborhood around $q_1 = 1$, given that $\rho < 1$), then $\beta^{sg} > \beta^{mg}$, which means that when the market thinks the manager is playing the growth strategy, it puts more weight on sales than on margins. These inequalities suggest the possibility of multiple equilibria in which the manager may cater to the market's beliefs. For example, if the market believes that the manager is playing a margins strategy, the pricing rule will tend to reward strong performance on the margins dimension more, which will, all else equal, make the margins strategy more attractive to the manager.

C.2. The Potential for Multiple Equilibria

We are now ready to construct the equilibria of the game between the manager and the market. This involves (i) postulating a market conjecture, $e^* \in \{0, 1\}$; (ii) then, using the updating rules derived above, computing the expected payoff to the manager under alternative effort allocations; and (iii) from there deriving the necessary and sufficient conditions under which $e = e^*$ is indeed the optimal strategy for the firm manager.

Suppose first that the market conjectures that the manager is following the growth strategy $e^* = 1$. Then, if the manager chooses to be consistent with the market conjecture by pursuing the growth strategy ($e = 1$), her expected sales, profits, and margins are respectively equal to

$$E(s_1 | e = 1) = Aq_1 = E(\pi_1 | e = 1); E(m_1 | e = 1) = 0. \quad (15)$$

From our analysis in the previous subsection, her expectation of the resulting stock price is therefore equal to

$$E(P_1^g | e = 1) = \delta(A + \beta^{sg} E(s_1 - Aq_1 | e = 1) + \beta^{mg} E(m_1 | e = 1)) = \delta A. \quad (16)$$

Thus, by choosing a growth strategy in the face of a growth conjecture, the manager obtains a total expected utility payoff of

$$E(\pi_1 + \alpha P_1^g | e = 1) = Aq_1 + \alpha \delta A. \quad (17)$$

Now, if the manager instead chooses to go against the market conjecture by focusing on margins ($e = 0$), her expected margins, profits, and sales are equal to

$$E(m_1 | e = 0) = A = E(\pi_1 | e = 0); E(s_1 | e = 0) = 0, \quad (18)$$

and her expectation of the resulting stock price is equal to

$$\begin{aligned} E(P_1^g | e = 0) &= \delta(A + \beta^{sg} E(s_1 - Aq_1 | e = 0) + \beta^{mg} E(m_1 | e = 0)) \\ &= \delta(A - \beta^{sg} Aq_1 + \beta^{mg} A). \end{aligned} \quad (19)$$

Thus, by choosing a margins strategy in the face of a growth conjecture, the manager obtains a total expected utility payoff equal to

$$E(\pi_1 + \alpha P_1^g | e = 0) = A + \alpha\delta(A - \beta^{sg} Aq_1 + \beta^{mg} A). \quad (20)$$

Comparing equations (17) and (20), we see that by deviating away from the market conjecture, the manager generates sales that are below expectations by an average amount Aq_1 , and margins that are above expectations by an average amount A ; the former translates into a loss in reputation and hence in the stock price, while the latter translates into a gain. Overall, the manager will choose to fulfill the market's conjecture by choosing $e = 1$ whenever

$$q_1 \geq 1 - \alpha\delta(\beta^{sg} q_1 - \beta^{mg}). \quad (21)$$

Thus, for sufficiently large market capacity q_1 , it will always be an equilibrium for the manager to play the growth strategy. Denote by q^- the value of q_1 that satisfies the above constraint with equality. Using the expressions for β^{sg} and β^{mg} derived above, one can analyze how the cutoff point q^- varies with the parameters of the model. In particular, note that q^- solves

$$q - \alpha\delta\rho = 1 - \alpha\delta \frac{q^2 v^a + \rho q v^s}{q^2 v^a + v^s}. \quad (22)$$

Using this fact, one can show that q^- is a decreasing function of both ability uncertainty, v^a , and the intensity of the manager's concern for the stock price, α . If either of these parameters goes to zero, it is easy to see that q^- goes to one. That is, we converge to the first-best situation, where the growth strategy is only played if $q_1 \geq 1$.

Applying similar reasoning, we can derive the conditions under which the margins strategy is an equilibrium. Suppose the market conjectures that the manager is playing a margins strategy. Then the expected utility to the manager from going along with the market and playing the margins strategy is equal to

$$E(\pi_1 + \alpha P_1^m | e = 0) = A + \alpha\delta A. \quad (23)$$

On the other hand, if the manager deviates from the market conjecture and chooses a growth strategy, her expected utility is equal to

$$E(\pi_1 + \alpha P_1^m | e = 1) = Aq_1 + \alpha\delta(A + \beta^{sm} Aq_1 - \beta^{mm} A). \quad (24)$$

By deviating from the market's conjecture, the manager reduces margins below expectations by an average amount A , whereas she raises sales above expectations by an average amount Aq_1 . The manager will choose to conform with the market's expectation and emphasize margins whenever

$$q_1 \leq 1 + \alpha\delta(\beta^{mm} - \beta^{sm} q_1). \quad (25)$$

Denote by q^+ the value of q_1 for which condition (25) holds with equality. Or, equivalently, call q^+ the solution to

$$q + \alpha \delta \rho q = 1 + \alpha \delta \frac{v^a + \rho v^m}{v^a + v^m}. \quad (26)$$

Once again, it is straightforward to show that q^+ is an increasing function of both v^a and α . Moreover, we have

$$q^+ > 1 > q^- \quad (27)$$

when v^a and α are both strictly positive, whereas

$$q^+ = 1 = q^- \quad (28)$$

in the first-best case when at least one of v^a and α is equal to zero.

Our analysis of the static model is summarized in the following proposition:

PROPOSITION 1: *Suppose that α and v^a are both strictly positive. Then there exist two cutoff values q^- and q^+ such that: (i) $q^+ > 1 > q^-$; (ii) if $q_1 < q^-$, the unique equilibrium is that the firm plays the margins strategy ($e = 0$); (iii) if $q_1 > q^+$, the unique equilibrium is that the firm plays the growth strategy ($e = 1$); and (iv) if $q^- \leq q_1 \leq q^+$, there are two equilibria, one where the firm plays margins and the other where the firm plays growth. The size of this multiple-equilibrium region increases with α or v^a , but it vanishes as α or v^a converge to zero, in which case the firm always plays the first-best strategy.*

Again, the economic intuition behind the proposition is one of managers catering to the stock market's beliefs. If the market perceives that a firm is trying hard to generate sales growth, it will tend to react more strongly to news about growth because such news is more informative about managerial ability. In contrast, if the market thinks that the firm is focusing its efforts on improving margins, it will tend to react more strongly to news about profitability. In either case, a manager who is concerned about stock prices will tend to give the market what it is looking for, which creates the scope for multiple equilibria. And the more intense the manager's concern with stock prices, the wider the range of parameters under which multiple equilibria arise.

II. Dynamics

In the static analysis, we took the market size q_1 to be exogenously fixed. Now we consider the case in which q varies endogenously over time in response to firms' past strategic choices. In particular, we assume that the longer a firm pursues a growth strategy, the more it exhausts the potential pool of customers for its products, and hence the lower is demand going forward. This formulation yields a system with its own internal dynamics, and allows us to ask, for example, how the volatility of sales growth depends on the various parameters of the model.

A. The Evolution of Demand

Imagine an infinite-horizon economy in which, at each date t , q_t denotes the per-firm unfulfilled demand for a durable good that is produced by many

identical firms. This demand comes from a unit mass of consumers with unfilled orders. Every period, a fraction γ of these consumers die and are replaced by new consumers with individual orders of Q_t . These births and deaths lead to different dynamic equations for q_t depending upon whether firm managers choose a growth versus a margins strategy. More specifically, the dynamics of unfilled orders q_t can be described as follows.

In the growth regime, there are per-firm sales of Aq_t , leaving unfilled orders of $(1 - A)q_t$ at the end of period t .¹⁰ Before the next period, a fraction γ of consumers with unfilled orders die, so the carryover into period $t + 1$ is $(1 - \gamma)(1 - A)q_t$. In addition, a fraction γ of new consumers are born, with orders of Q_t . So if the growth regime prevails at time t , we have

$$q_{t+1} = (1 - \gamma)(1 - A)q_t + \gamma Q_t. \quad (29)$$

In the margins regime, everything is the same as above except that per-firm sales at time t are zero. So if the margins regime prevails at time t , we have

$$q_{t+1} = (1 - \gamma)q_t + \gamma Q_t. \quad (30)$$

More generally, if managers choose to allocate an interior level of effort e_t to the growth strategy—so that we are in neither the pure growth nor the pure margins regime—we have

$$q_{t+1} = (1 - \gamma)(1 - Ae_t)q_t + \gamma Q_t. \quad (31)$$

These equations fully characterize the evolution of demand from one period to the next. Again, the key point to note is that if we are in the growth regime at time t , this leads to lower demand at time $t + 1$, since the representative firm produces more and hence satisfies a greater fraction of the unfilled orders.

B. Information Structure

With respect to incentives and information structure, we try to stick as closely as possible to the assumptions of the static model. Firms are run by a succession of one-period managers. The manager who is in charge at time t has ability a_t , and a utility function given by

$$U_t = \pi_t + \alpha P_t. \quad (32)$$

Here, π_t denotes time- t profits, and P_t is the time- t market value of the firm. That is,

$$P_t = E_t \sum_{k \geq 1} \delta^k \pi_{t+k}, \quad (33)$$

¹⁰ We assume that there are many identical firms selling to the same pool of customers, so that we can ignore the effect of firm-level shocks to sales and managerial ability on the aggregate dynamics of unfilled orders. In other words, these firm-level shocks wash out across the many firms in the industry.

where, again, δ denotes the market's discount factor from one period to the next.

By analogy to the static model, we assume that the ability of the time- t manager affects both current and future profits, which in turn implies that it also matters for the firm's time- t stock market value P_t . This assumption corresponds to the reasonable idea that a better manager leaves a firm in better shape when she departs, all else equal, that is, the benefits of her tenure show up in profits even after she is gone. In particular, profits are given by

$$\pi_t = a_t e_t q_t + a_t(1 - e_t) + \varepsilon_t^s + \varepsilon_t^m + a_{t-1} + \rho(\varepsilon_{t-1}^s + \varepsilon_{t-1}^m), \quad (34)$$

where ε_t^s and ε_t^m denote the time- t realizations of the sales and margin shocks. Note that in addition to the carryover of previous-manager ability, sales and margin shocks continue to be partially persistent. As before, we continue to assume that profits can be decomposed into a sales and a margin component, each of which can be separately observed:¹¹

$$\pi_t = s_t + m_t, \quad (35)$$

where

$$s_t = a_t e_t q_t + \varepsilon_t^s + \rho \varepsilon_{t-1}^s \quad (36)$$

and

$$m_t = a_t(1 - e_t) + a_{t-1} + \varepsilon_t^m + \rho \varepsilon_{t-1}^m. \quad (37)$$

All random variables are taken to be independent and serially uncorrelated. Moreover, to simplify the analysis, we assume that at the beginning of period t , the past shocks ε_{t-1}^s , and ε_{t-1}^m , as well as the ability a_{t-1} of the previous manager, are all publicly revealed. This is a modeling device that makes the task of forecasting future profits exactly the same as in the static case. In particular, the market can now make forecasts based on "adjusted values" of sales and margins, which factor out a_{t-1} , ε_{t-1}^s , and ε_{t-1}^m . That is, the market can create $s_t^a = a_t e_t q_t + \varepsilon_t^s$ and $m_t^a = a_t(1 - e_t) + \varepsilon_t^m$, and use these two variables to predict next period's profits.

The timing of events can be described as follows. First, at the beginning of every period t , the new manager chooses how to allocate her effort; this, together with past managerial ability and the realization of current and past shocks, determines the level of current profits, sales, and margins. Second, based on observed sales and margins, and its own prior conjecture as to the strategy that the firm is following, the market generates a forecast of next period's profits, which in turn influences the stock price P_t . Third, starting from a given amount q_t of unfilled orders at the beginning of period t , firms' strategic choices determine the level of unfilled orders q_{t+1} at the beginning of the subsequent period $t + 1$. Finally, at the beginning of period $t + 1$, past

¹¹ In this decomposition, we assume that the ability a_{t-1} of the previous period's manager influences this period's margins, as opposed to sales.

shocks and past managerial abilities become public information and the whole cycle begins again.

We shall focus attention on *dynamic self-fulfilling equilibria*, which are defined as sequences of unfilled orders q_t , market conjectures e_t^* , and managerial effort allocations e_t , such that, for all t : (i) given (q_t, e_t^*) it is optimal for managers to choose effort allocation e_t ; (ii) the allocation e_t coincides with the market conjecture e_t^* ; and (iii) q_{t+1} satisfies the general version of the transition equation (31).

C. First-Best Action

Before proceeding, we consider the first-best benchmark in two cases. In the first of these, the demand flow Q_t remains constant over time. In the second, Q_t declines monotonically.

C.1. A Market with Constant Demand Flow

Suppose that the demand flow Q_t is simply equal to a constant Q at all times, and that either α or ν^a are equal to zero, so that managers do not care about the market. We know from our analysis of the static model that when α or ν^a are equal to zero, managers will be indifferent between the growth and margins strategies at date t if and only if $q_t = 1$. In the dynamic case, the question is whether we can sustain a steady-state equilibrium in which $q_t = 1$ for all t , and in which managers devote a fixed interior level of effort $e^f \in (0, 1)$ to growth.

If the manager's effort allocation is indeed e^f , equation (31) tells us that the dynamics of q_t are given by

$$q_{t+1} = (1 - \gamma)(1 - Ae^f)q_t + \gamma Q. \quad (38)$$

This implies a unique steady-state value of q_t , namely,

$$q^s = \gamma Q / (\gamma + (1 - \gamma)Ae^f). \quad (39)$$

Therefore, a necessary and sufficient condition for the interior amount of effort to be chosen in all periods is that q^s be equal to one. This in turn allows us to solve for the first-best effort allocation e^f , namely,

$$e^f = \gamma(Q - 1) / ((1 - \gamma)A). \quad (40)$$

It follows that in the first-best case where managers do not care about impressing the market, a steady-state equilibrium—in which growth-related effort is given by e^f every period and $q_t \equiv 1$ —always exists as long as the parameters γ , Q , and A are such that the right-hand side of (40) lies strictly between zero and one. Note also that this equilibrium is stable in the sense that starting from $q_t > 1$, managers will choose the growth strategy so that q_t will decrease towards one. Similarly, starting from $q_t < 1$, managers will choose the margins strategy which will make q_t increase towards one. In contrast, we will show

below that when managers do cater to the market, it is generally not possible to sustain a stable steady-state equilibrium of this sort.

If the parameters are such that the right-hand side of (40) does not lie between zero and one, we have a corner solution. In particular, if the right-hand side of (40) exceeds one, this means that the demand flow is so great that even under the pure growth strategy, the firm cannot keep up with demand and unfilled orders asymptote to $q_t = \gamma Q / (\gamma + (1 - \gamma)A) > 1$. Similarly, if $Q < 1$, so that the right-hand side of (40) is negative, the firm plays the pure margins strategy and unfilled orders asymptote to Q .

C.2. A Market with Shrinking Demand Flow

It is straightforward to extend the above reasoning to the case in which Q_t declines monotonically. Consider the natural analog to (40) with a time-varying effort e_t^f :

$$e_t^f = \gamma(Q_t - 1) / ((1 - \gamma)A). \quad (41)$$

If we start out with a value of Q_t such that the right-hand side of (41) exceeds one, the firm is initially at a corner solution and plays the all-growth strategy in which $e_t^f = 1$. Over time, as Q_t falls, the right-hand side of (41) drops below one, at which point we have an interior solution with the level of effort devoted to growth, e_t^f , given by (41). From this point on, the equation tells us that further declines in Q_t lead to matching declines in e_t^f . Eventually, when Q_t drops below one, we hit the other corner, where $e_t^f = 0$, and the firm plays the all-margins strategy. The important point here is that if demand shrinks monotonically, so must the effort devoted to the growth strategy. As we will demonstrate below, this need not be the case when managers care about stock prices.

D. Excess Volatility When Managers Cater to the Market

We now reintroduce a managerial concern with stock prices, and show that this can lead to endogenous volatility in sales growth. Our reasoning proceeds in two steps. First, we consider a given period t and show that Proposition 1 continues to hold. That is, for the same cutoff values q^- and q^+ as in the static case: (i) the unique equilibrium is that firms choose the growth strategy whenever $q_t > q^+$; (ii) the unique equilibrium is that firms choose the margins strategy whenever $q_t < q^-$; and (iii) for $q^- < q_t < q^+$, the two equilibria coexist.

Second, we consider the dynamics of q_t and the resulting managerial decisions over time. To do so, we assume that in a dynamic equilibrium the market maintains the same conjecture as in the previous period whenever there continues to exist a static equilibrium with that conjecture. Thus, for example, if q_t enters the multiple-equilibrium region (q^-, q^+) from above—that is, coming from the region $q_t > q^+$ where the growth strategy is the unique static equilibrium—then as long as q_t remains in that region, the firm and the market keep playing the

growth equilibrium. Or put differently, whenever there is a choice of equilibria, we select the one played in the prior period.

D.1. Within-Period Analysis

To understand why the same cutoff values q^- and q^+ obtain as in the static model, observe that a firm's market value as of the end of period t can be expressed as

$$P_t = \delta E(a_t + \rho(\varepsilon_t^s + \varepsilon_t^m) + \Omega_{t+1} \mid s_t^a, m_t^a, e_t^*), \quad (42)$$

where Ω_{t+1} is a function that depends on only those random variables that are realized from $t + 1$ onwards. Thus, as of period t , the market makes a conjecture e_t^* about the firm's current choice of strategy, and tries to forecast exactly the same item as in the static model, namely, $\delta(a_t + \rho(\varepsilon_t^s + \varepsilon_t^m))$; period- t observations of profits, sales, and margins are of no help in forecasting Ω_{t+1} .¹² Moreover, given our previous assumptions on information structure, the market has access to the same two predictive variables as before, $s_t^a = a_t e_t q_t + \varepsilon_t^s$ and $m_t^a = a_t(1 - e_t) + \varepsilon_t^m$. So the inference problem is completely unchanged.

Similarly, from the perspective of the period- t manager, everything is also the same, since her utility function is the same as in the static model. Thus, Proposition 1 continues to apply as stated, with the only modification being that q_1 is replaced everywhere with q_t .

D.2. Equilibrium Selection in the Dynamic Model

In order to fully pin down the dynamic evolution of q_t and e_t , we need to specify a mechanism for picking an equilibrium when we are in the parameter region (q^-, q^+) where there are multiple static equilibria. To do so, we make the following assumption:

ASSUMPTION 1: *If at date t , for a given value of q_t there exist (according to Proposition 1) multiple static equilibria, the same equilibrium is chosen as in the previous period.*

As noted above, this assumption can be thought of as introducing an element of history dependence into the equilibrium selection criterion.¹³ Importantly, the assumption does not require any departure from Bayesian rationality on the part of either the market or managers. Rather, it follows from imposing some relatively mild noncommon-knowledge structure on the market's higher-order beliefs out of equilibrium. In particular, Assumption 1 would be implied by the following higher-order beliefs. Suppose that the market believes that managers believe that the market is inertial in its conjectures—in other words,

¹² In addition to future realizations of the random variables, Ω_{t+1} also depends on the future path of unfilled orders q_{t+1} , q_{t+2} , etc., which in turn shape future firm strategies e_{t+1} , e_{t+2} , etc. But as we will show below, these variables evolve completely deterministically in equilibrium.

¹³ Rajan (1994) makes an identical assumption in a model with a similar dynamic structure.

the market believes that managers believe that the market always maintains the same conjecture about managers' strategies as in the prior period. And suppose further that managers understand that the market has such beliefs about them. This will lead both sides to coordinate on the equilibrium selection criterion we have assumed. Moreover, neither side's beliefs will ever be contradicted along the path of play.

D.3. Deterministic Cycles with Constant Demand Flow

With all the pieces now in place, let us return to the case in which the demand flow Q_t is equal to a constant Q . We can now state our central result:

PROPOSITION 2: *Under Assumption 1, with Q_t constant, and with α and v^a both positive, we have that: (i) there no longer exists a stable stationary equilibrium in which managers allocate a positive and time-invariant amount of effort to both growth and margins, except for possibly on a zero-measure subset of parameters; and (ii) there exists a subset of parameters with positive measure for which one can construct a cyclical equilibrium. In this cyclical equilibrium, the firm switches back and forth between the growth and margins regimes at regular intervals, and the level of unfilled orders q_t fluctuates within a range bounded by two extreme values, $\bar{q} > q^+$ and $\underline{q} < q^-$.*

Proof: Part (i) of Proposition 2 is proven in Appendix B. Here, we just establish part (ii); this part of the proof is more brief and captures most of the important intuition. Let q^g and q^m denote the fixed points of the dynamic equations (29) and (30), that is, the steady-state values of q_t in the all-growth and all-margins regimes, respectively. We have already noted that

$$q^g = \gamma Q / (\gamma + (1 - \gamma)A) \quad (43)$$

and:

$$q^m = Q. \quad (44)$$

Then, a sufficient condition for the existence of endogenous cycles is simply that

$$q^g < q^- \text{ and } q^m > q^+. \quad (45)$$

To see why, suppose that at some date t , unfilled orders q_t lie in the interval (q^g, q^-) . Then, the unique static equilibrium is for managers to play the margins strategy. But then Assumption 1 implies that the market will maintain the margins conjecture, and managers will continue to fulfill that conjecture, until q_t crosses the threshold q^+ . This crossing will necessarily occur in finite time, since $q^m > q^+$. When the threshold is crossed, the margins strategy will cease to be an equilibrium and the firm will switch to the growth strategy. Once the growth strategy is in place, q_t will begin to decline, heading towards q^g . Now, since $q^g < q^-$, the sequence q_t will cross the threshold q^- in finite time, which




Figure 1. Deterministic cycles with constant demand flow. The parameter values are chosen as follows: $Q = 2$, $\gamma = 0.1$, $A = 0.5$, $v^a = v^s = v^m = 1$, $\rho = 0.1$, and $\delta = 1$. We then plot the evolution of q_t for three cases: (i) $\alpha = 0$, (ii) $\alpha = 1.0$, and (iii) $\alpha = 1.5$.

in turn brings us back to the situation in which we started, that is, q_t belongs to the interval (q^g, q^-) and the margins strategy is the unique equilibrium.

This establishes part (ii) of the proposition. It also should be pointed out that, since q^- and q^+ are independent of Q and A , the above condition (45) is easily satisfied if γ is small relative to A and Q because in this case, q^g is approximately equal to $\gamma Q/A$, which can be made arbitrarily small, while $q^m = Q$, which can independently be made large. Q.E.D.

Figure 1 provides an illustration of deterministic cycles with constant demand flow. We choose the following baseline parameter values: $Q = 2$, $\gamma = 0.1$, $A = 0.5$, $v^a = v^s = v^m = 1$, $\rho = 0.1$, and $\delta = 1$. We then solve the model forward over 100 periods, for each of three values of α : (i) $\alpha = 0$ (i.e., the first-best case); (ii) $\alpha = 1.0$; and (iii) $\alpha = 1.5$. In each case, we plot the evolution of q_t over time.

In the first-best case, q_t starts out at a value of 0.2, so we are initially in the margins regime. The low level of output associated with the margins strategy causes q_t to rise. When q_t hits one, the firm switches to an interior level of effort given by (40). From this point onward, the firm is in steady state, with $e^f = 0.22$ and q_t fixed at one.

Moving to the catering-to-the-market cases, when $\alpha = 1.0$ we have that $q^- = 0.71$ and $q^+ = 1.41$. Thus, the firm stays with the margins strategy—and q_t keeps increasing—longer than in the first-best case, until q_t crosses 1.41. At this point, the firm switches to the growth strategy, which drives q_t sharply

downward. When q_t falls below 0.71, the firm goes back to the margins strategy and the cycle starts over again.

Things are qualitatively similar when $\alpha = 1.5$. However, we now have that $q^- = 0.64$ and $q^+ = 1.59$, that is, the size of the multiple-equilibrium range has increased. This leads to more pronounced overshooting, and hence to a smaller number of cycles, with each cycle being of larger amplitude. Consider, in particular, the first switch point. With $\alpha = 1.5$, the firm stays with the margins strategy longer, until q_t crosses 1.59. Thus, the first transition to the growth strategy is delayed, coming at $t = 14$, rather than at $t = 11$, as in the case where $\alpha = 1.0$. The bottom line is that the increase in α results in q_t being on average further away from the first-best value of one. The natural interpretation is that a more intense managerial focus on current stock prices has negative consequences for economic efficiency.

D.4. Nonmonotonic Adjustment with Shrinking Demand Flow

The other case of interest is the one in which the demand flow declines monotonically over time. As seen above, in the first-best case this decline in demand flow is matched by a monotonic reduction in the effort devoted to the growth strategy. The goal here is to show that when managers cater to the stock market, the adjustment of growth-related effort—and hence of q_t —can be nonmonotonic.

The mechanism here is very similar to the mechanism that generates cycles in the constant- Q case. For the sake of concreteness, suppose that the demand flow Q_t declines geometrically over time, that is

$$Q_t = Q_0 \beta^t, \quad (46)$$

with $0 < \beta < 1$. Moreover, let Q_0 be large, so that the firm starts out in the growth regime. In this regime, the dynamics of q_t are given by

$$q_{t+1} = (1 - \gamma)(1 - A)q_t + \gamma Q_0 \beta^t. \quad (47)$$

According to (47), q_t will eventually start declining. Moreover, by Assumption 1, the firm will stay in the growth regime past the point where q_t drops below one, until q_t passes through q^- . (Recall, by contrast, that in the first-best case, the firm transitions to an interior level of effort as soon as q_t reaches one.) Once q^- has been breached, the firm will switch to the margins regime, where the dynamics of q_t are given by

$$q_{t+1} = (1 - \gamma)q_t + \gamma Q_0 \beta^t. \quad (48)$$

Because the effort devoted to growth has been cut by a discrete amount, it is possible that q_t will now start rising, in spite of the general downward drag from the shrinking demand flow. Indeed, q_t may well cross back above q^+ , leading the firm to switch back to the growth regime once again.

Of course, in the long run, q_t will always converge to zero (since the demand flow Q_t is going to zero), and the firm will wind up absorbed in the margins

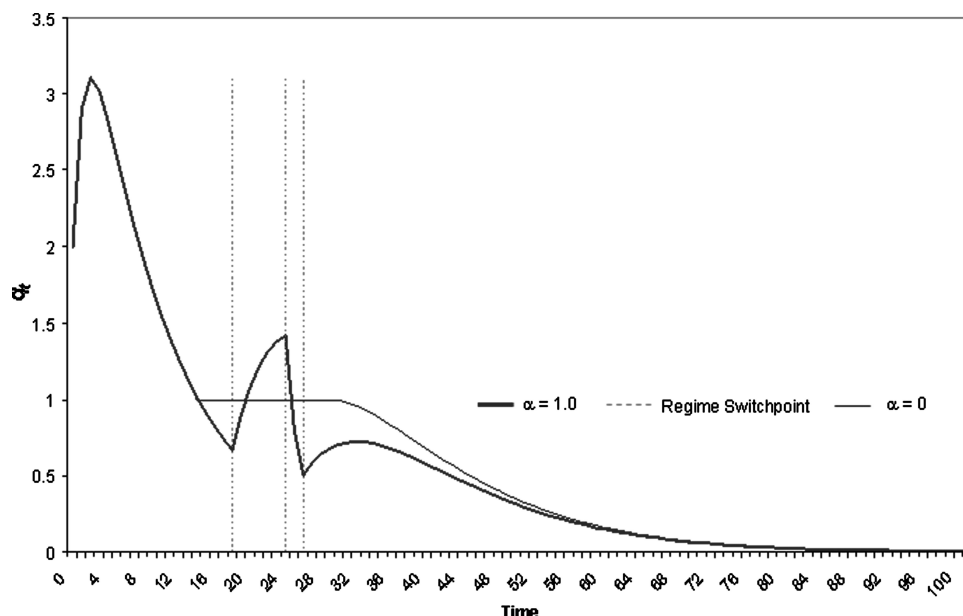


Figure 2. Nonmonotonic adjustment with shrinking demand flow. The parameter values are chosen as follows: $Q_0 = 20$, $\beta = 0.9$, $\gamma = 0.1$, $A = 0.5$, $v^a = v^s = v^m = 1$, $\rho = 0.1$, and $\delta = 1$. We then plot the evolution of q_t for two cases: (i) $\alpha = 0$ and (ii) $\alpha = 1.0$.

regime. But what is interesting relative to the first-best case is that this convergence need not be monotonic. Indeed, in principle it can involve a number of transitions back and forth between the growth and margin regimes, with corresponding fluctuations in the value of q_t .

Figure 2 illustrates this phenomenon of nonmonotonic convergence. We set $Q_0 = 20$, $\beta = 0.9$, $\gamma = 0.1$, $A = 0.5$, $v^a = v^s = v^m = 1$, $\rho = 0.1$, and $\delta = 1$. We then plot the evolution of q_t in two cases: (i) the first-best case where $\alpha = 0$; and (ii) $\alpha = 1.0$.

In the first-best case, q_t starts out at a value of 2.0, so we are initially in the growth regime. In spite of this fact, the demand flow is so high in the early periods that q_t continues to increase, reaching a peak value of 3.2 at $t = 3$. From this point on, the level of output associated with the growth strategy is enough to overwhelm the now-smaller demand flow, so q_t begins to fall. When q_t hits one, which happens at $t = 14$, the firm switches to an interior level of effort given by (41). This effort level declines gradually over time, and q_t remains fixed at one, until $t = 30$. From this point onward, the demand flow is too low for an interior effort level to be viable. The firm therefore moves to the pure margins strategy and q_t declines toward zero.

In the second-best case, with $\alpha = 1.0$, we have as before that $q^- = 0.71$ and $q^+ = 1.41$. Everything works the same as in the first-best case during the early periods—the firm plays the growth strategy, and q_t first rises, then begins to fall. However, the two cases diverge when q_t reaches one. Rather than moving

to an interior effort level, the firm stays with the growth strategy longer, until q_t drops below 0.71, which happens when $t = 18$. At this point, the firm switches to the margins strategy. Demand flow is still strong enough that this switch leads q_t to begin increasing. When q_t rises above 1.41, the firm goes back to the growth strategy and q_t starts falling again. As it turns out, this second growth episode is the last one; the next time that the firm returns to the margins strategy, demand flow has abated to the point where q_t never again reaches q^+ . Thus, overall, the firm starts with growth, switches to margins, and goes back to growth, all before settling permanently into the margins strategy.

D.5. Possible Mechanisms for Restoring Efficient Behavior

Readers have suggested to us a couple of mechanisms that might potentially eliminate the inefficient excess volatility in our model. The first is to base equilibrium selection not on the history dependence criterion embodied in Assumption 1, but rather on managerial cheap talk. In particular, if at date t there exist multiple equilibria, one might imagine that the incumbent manager simply announces which strategy she intends to pursue, and the market believes her. At first glance, it might appear that this mechanism would allow for an approximately efficient outcome whereby the growth strategy is played whenever $q_t > 1$, and the margins strategy is played whenever $q_t < 1$.

However, this conclusion does not hold if we perturb the model slightly so that the manager has private information about the upcoming realizations of the sales and margins shocks, ε_t^s and ε_t^m . (This perturbation does not affect our results if we continue to disregard announcements by the manager.) To see why, suppose that the market believes the manager's announcements, and that within each equilibrium, the updating rules continue to be given as before. Now consider a situation where $q_t < 1$ —so that it is efficient for the manager to announce the margins strategy—but where the manager anticipates a large positive sales shock. The manager will prefer to announce the growth strategy because she wants to get more credit for a strong showing on the sales dimension. Thus, efficient equilibrium selection is not attainable via a mechanism that relies on managerial announcements because of the opportunistic nature of these announcements with respect to managers' private information.

Even if managerial cheap talk is not helpful, it is possible that the firm might try to more credibly signal its upcoming strategy via the choice of a manager with a specific set of skills (e.g., a cost-cutting specialist as opposed to a growth specialist). We certainly would not want to take the position that managerial changes cannot be an important signal about firm strategy in some cases. Nevertheless, this type of signaling is unlikely to completely drive out the effects that we are interested in. One countervailing force is that because the signal is by its nature costly, it will not always be optimal to use it. To take one example, suppose a firm that is currently pursuing a growth strategy has an incumbent manager with a great deal of firm-specific expertise. If the firm wants to signal a deviation to a cost-cutting strategy, it would have to hire an outsider with a reputation as a cost-cutter. But in spite of this reputation, the outsider

may simply not have as much firm-specific human capital as the incumbent generalist, and hence it would be efficiency-reducing to change managers.

Furthermore, even when managerial changes do happen and are informative, they will not generally restore the first-best outcome without cycles. Suppose that the first-best steady state calls for some interior value of effort devoted to the growth strategy—e.g., in the example illustrated in Figure 1, we had $e^f = 0.22$. Unless it is possible to find a manager with an ironclad reputation for choosing exactly $e = 0.22$ and never wavering from this policy (say to $e = 0.40$), it will be impossible to sustain the first-best, and all of our basic results will continue to go through.

Intuitively, one might think of managerial reputation as shrinking the support for e , rather than as communicating to the market a single precise value of e . Thus, instead of believing that e can take on any value between zero and one, if we have a manager with a reputation for cost-cutting, the market might think that e can only vary between zero and $\frac{1}{2}$. But as long as the precise value of e chosen by the manager remains contemporaneously unobservable within this truncated support, our qualitative results will continue to apply, even if the quantitative magnitudes are tempered somewhat.

E. Behavioral Variations

E.1. Investors Use Simple Models

So far we have taken stock market investors to be fully rational, that is, able to form Bayesian forecasts based on all observed variables. In what follows, we briefly explore the consequences of a simple form of bounded rationality. In particular, we suppose that there is a representative investor who can only run univariate regressions.¹⁴ Since optimal forecasting in our setting requires the use of bivariate regressions, this is a meaningful limitation. In all other respects, however, the representative investor is completely rational. His univariate forecasts are unbiased, and in any given equilibrium (growth or margins) he picks the *single variable* that minimizes the variance of forecast errors.

In principle, the single most powerful forecasting variable for next period's profits can be current profits, current sales, or current margins, depending on the equilibrium that we are in. So to solve the general case, we would have to compute the forecast errors associated with each of these three variables, across both the growth and margins regimes. To keep things simple, we focus on an extreme but illustrative case, where the persistence coefficient ρ is small but nonzero. We first revisit the static model, and then sketch the dynamic implications.

In the context of the static model, it is easy to see that when ρ is positive but small, the best predictor of π_2 in the growth regime is s_1 , and the best predictor

¹⁴ This assumption can be motivated by the observation that investors tend to simplify the world around them, and hence do not use all available sources of information in making their forecasts. See Hong and Stein (2003) for an extensive discussion and references to the relevant psychology literature.

of π_2 in the margins regime is m_1 . Moreover, since s_1 and m_1 are uncorrelated, the optimal univariate regression coefficients β^{sg} and β^{mm} are exactly the same as given before. All that changes is that β^{mg} and β^{sm} are now effectively set to zero, as opposed to their previous values of ρ in the fully rational case. In other words, in the growth regime the representative investor focuses only on sales, and gives it the same weight as before, while completely ignoring margins, with the reverse being true in the margins regime.

As in Section I, we can calculate cutoff values of q_1 that define the regions over which the growth and margins equilibria can exist. Denoting these revised cutoffs by q_b^- and q_b^+ , we have that q_b^- solves

$$q = 1 - \alpha\delta \frac{q^2 v^a + \rho q v^s}{q^2 v^a + v^s}. \quad (49)$$

This equation is almost identical to its counterpart (22) in the rational case, except that the term $(-\alpha\delta\rho)$ on the left-hand side of that equation has now disappeared. It follows immediately that

$$q_b^- < q^-. \quad (50)$$

Similarly, q_b^+ satisfies

$$q = 1 + \alpha\delta \frac{v^a + \rho v^m}{v^a + v^m}. \quad (51)$$

This equation also looks like its counterpart (26) in the rational case, except for the term $(\alpha\delta\rho q)$, which has disappeared from the lefthand side. This in turn yields

$$q_b^+ > q^+. \quad (52)$$

Thus, for the static model, introducing the bounded-rationality feature enlarges the multiple-equilibrium region from (q^-, q^+) to (q_b^-, q_b^+) . There is a very natural intuition for this result. Suppose we are currently in a growth equilibrium. With bounded rationality, the market pays attention to sales as before, but now completely ignores margins. Thus, for a manager who contemplates deviating from the growth strategy to the margins strategy, the prospective hit to her stock price is more severe than before—the market will penalize her for the reduction in sales, but will give her no credit whatsoever for the improvement in margins. As a result, she is more inclined to simply continue playing the growth strategy. Similarly, if we start off in a margins equilibrium, the manager will again be very reluctant to deviate to a growth strategy.

Notice that introducing bounded rationality has an effect that is isomorphic to increasing the parameter α , which measures how intensely managers care about the current stock price. In the former case, the potential for multiple equilibria goes up because there is a bigger stock price penalty for bucking the expectations of a naive market; in the latter case, the potential for multiple equilibria goes up because managers are less willing to tolerate a stock price penalty of any given size.




Figure 3. The effect of bounded rationality with constant demand flow. The parameter values are chosen as follows: $Q = 2$, $\gamma = 0.1$, $A = 0.5$, $v^a = v^s = v^m = 1$, $\rho = 0.1$, and $\delta = 1$. We then plot the evolution of q_t for two cases: (i) the case in which investors are fully rational, and (ii) the case in which investors are boundedly rational.

When we move to the dynamic setting, everything continues to work exactly as before, except that we now use the new cutoff values q_b^- and q_b^+ in place of their rational counterparts q^- and q^+ . This implies that whenever the conditions of Proposition 2 are satisfied and a cyclical equilibrium exists, the magnitude of the fluctuations is greater than before, in much the same way as an increase in α leads to more extreme fluctuations. Again, this is because once entrenched in a given regime, managers are more reluctant than before to change strategies. As a result, they play any given strategy longer, which leads to more pronounced overshooting of q_t .

Figure 3 illustrates the contrast between the rational and bounded-rationality cases. We return to a constant-demand-flow environment and set $Q = 2$, $\gamma = 0.1$, $A = 0.5$, $v^a = v^s = v^m = 1$, $\rho = 0.1$, $\delta = 1$, and $\alpha = 1$. We then trace out the evolution of q_t for both cases. The plot corresponding to the fully rational case is actually identical to the one for $\alpha = 1$ in Figure 1; all the other parameter values are the same, and so we again have $q^- = 0.71$ and $q^+ = 1.41$. With bounded rationality, $q_b^- = 0.65$ and $q_b^+ = 1.55$; these are almost the same cutoffs that we obtained under full rationality in Figure 1 by setting $\alpha = 1.5$. So not surprisingly, we get essentially the same outcome: cycles are of both longer duration and greater amplitude, and on average, q_t is farther away from the first-best value of one.

E.2. Investors Treat Stocks in the Same Category Similarly

Another potentially interesting extension of our model—which we do not pursue formally here—follows from the observation of Barberis and Shleifer

(2003) that investors tend to place stocks into broad categories and to think about stocks within a given category in a similar way (see also Mullainathan (2002)). In our context, this might mean using the same pricing rule (based on either a growth conjecture or a margins conjecture) to value all firms in a category at a given point in time. For example, it might be that during the dot-com boom, investors lumped all internet retailers into one category and used the same growth-oriented pricing rule to value all of them.

In such an environment, the behavior of any one firm can have interesting spillovers onto other firms in the same category. If Amazon.com initially has very strong growth prospects and winds up entrenched in the growth equilibrium, this might force other internet retailers into pursuing a growth strategy as well, because they know they will be evaluated by the market based on the pricing rule from the growth equilibrium.

III. Empirical Implications

We now turn to the model's empirical content. First, we discuss the implications that follow from the version with fully rational investors. Next, we briefly touch on the additional predictions that arise when bounded rationality is added to the mix.

A. The Model with Rational Investors

Unlike most of the literature on the stock market and real activity, our theory's primary focus is not on how the *level* of stock prices influences real variables. That is, we have little to say regarding the textbook *q*-theoretic prediction that a firm's investment and sales growth will be higher when its stock price is high (Tobin (1969)). Instead, what matters in our framework is the market's pricing rule, that is, the *sensitivity* of stock prices to observable variables such as sales and profit margins. Our central implication is thus that a firm's investment and sales growth will be higher when its stock price is more sensitive to growth-oriented metrics; and conversely, a firm's investment and sales growth will be lower—and profit margins higher—when its stock price is more sensitive to measures of costs and margins. Moreover, a subsidiary prediction is that these basic effects will be more pronounced in those firms where managers have high-powered incentives to maximize short-term stock prices.

How might one operationalize tests of these hypotheses? Although a complete treatment of this question is beyond the scope of the paper, we can offer a few tentative thoughts. First, one might begin by using cross-sectional regressions to estimate the sensitivity of stock prices to different variables at a given point in time. To be more specific, imagine that for industry *i* at time *t*, we take all the existing firms and run a cross-sectional regression of each firm's market-to-book ratio against: (i) some measure of its growth-related performance (e.g., sales growth over the prior year); and (ii) some measure of its per-unit profitability (e.g., its gross margin). The coefficient on the former variable would then be the "growth beta" for industry *i* at time *t*, and the coefficient on the latter variable would be the "margins beta."

The next step would be to relate changes in these growth and margins betas to changes in the nature of firm performance. In particular, one could ask whether, for example, investment and sales growth are on average higher for firms in industry i at those times when the industry's growth beta is larger. A finding along these lines would be broadly consistent with the predictions of our model.

At the same time, such a finding would not cut as cleanly as one would ideally like. In particular, it would also be consistent with a story whereby the firm's choice of strategy is always set at the first-best level, and the first-best strategy simply moves around in response to exogenously changing market or industry conditions.¹⁵ If the firm does shift its behavior over time in this exogenous fashion, one would expect the market's pricing rule to adapt accordingly—sometimes responding more to sales news, and sometimes responding more to margins news. However, this could potentially reflect only one direction of causality (from the firm's strategy to the market's pricing rule) and not the feedback effect that makes the firm's strategy endogenous, as envisioned by our model.

In order to get around these ambiguities, it may help to focus on a more sharply differentiating prediction of our theory: The link between firm strategy and the market's pricing rule ought to be more pronounced for firms with higher values of the short-termism parameter α . Intuitively, if $\alpha = 0$, strategies are always first-best in nature, and to the extent that they vary over time, this variation must reflect exogenous changes in market or industry conditions. However as α increases, we get the added catering-to-the-market feedback effect that is central to our model, and that should intensify the link between firm strategy and the market's pricing rule. In essence, what we are proposing is a differences-in-differences empirical specification in which we use firms with near-zero values of α (perhaps these are private firms, or firms in which managers hold large illiquid stakes) as a control group to address the above caveat about exogenously driven changes in firm strategy. Of course this differences-in-differences approach is complicated by the fact that observed variation in α may be partially endogenous, so ideally one would like to implement the tests using an instrument for α .

We are not aware of any empirical work that speaks to these predictions. However, we should mention a pair of recent papers by Baker and Wurgler (2004a,b) that share a similar spirit. Baker and Wurgler develop a catering theory of dividends and provide the following evidence in favor of it. First, they calculate at any point in time a "dividend premium," defined as the difference in the average market-to-book ratio of dividend-paying and nondividend-paying firms. They then show that variation over time in this dividend premium helps to explain firms' choices of dividend policy: Firms are significantly more likely to initiate dividend payments when the dividend premium is high. While Baker

¹⁵ For example, it may be optimal for a firm to focus entirely on growth in the early stages of its industry's development, so as to establish its product as the market standard or to otherwise lock up customers and suppliers. Later, once its position is established, it may be optimal for the firm to concentrate more on cost-cutting.

and Wurgler's interest is in a financial, rather than real variable (i.e., dividends, as opposed to sales growth or investment), the structure of their tests is closely analogous to what we propose above. That is, they estimate the point-in-time sensitivity of stock prices to a particular firm attribute using cross-sectional comparisons, and then show that managers' apparent willingness to deliver this attribute is greater when the market values it more highly.¹⁶

B. Consequences of Bounded Rationality

The fully rational version of the model has nothing to say about expected stock returns—they are simply constant over time. In contrast, the bounded rationality version of the model makes a novel and relatively subtle prediction. When we are in a growth equilibrium, the market fully impounds all growth-related information, but ignores margins-related information. This implies that in a growth equilibrium, firms with strong profit margins will be undervalued and have high expected returns, while firms with weak profit margins will be overvalued and have low expected returns. Conversely, in a margins equilibrium, the market takes account of margins-related information, but underweights growth-related information. So firms with strong growth performance will have high expected returns, and firms with weak growth performance will have low expected returns.

The key insight is that neither attribute—growth or margins—will be systematically mispriced at all times. Rather, there will be time variation in the nature of cross-sectional return predictability. For example, sometimes a trading strategy of buying high margin firms and shorting low margin ones will work well, and at other times it will work less well. Moreover, the theory tells us precisely when the strategy should be expected to work best: when the industry in question is in a high growth regime.

Again, we do not know of any work that can be thought of as a systematic test of this hypothesis. Perhaps the closest bit of evidence comes from Baker and Wurgler (2006). They too are interested in the general idea of time variation in the cross-section of expected stock returns, and they examine how the predictive content of different firm characteristics for stock returns varies with proxies for the overall level of investor sentiment. One of their findings is that: (i) when investor sentiment is relatively high, unprofitable firms have lower expected returns than profitable firms; but (ii) when investor sentiment is relatively low, this pattern reverses itself, so that unprofitable firms have higher expected returns. If one thinks of investor sentiment as being positively correlated with economy-wide growth prospects, this finding can be seen as fitting nicely with the predictions of our theory.

¹⁶ One distinction between our work and Baker and Wurgler (2004a, b) is that they interpret the time variation in the market's sensitivity to dividends as a manifestation of investor irrationality. In contrast, we have emphasized that time variation in the market's sensitivity to measures of growth and margins can be part of a fully rational equilibrium.

IV. Related Work

Our model relates to several distinct branches of literature. One of these, which we discussed briefly in the Introduction, considers the impact of investor sentiment and stock market mispricing on corporate investment. As noted above, our model differs from this work in that we do not require any irrationality on the part of investors to generate our principal results (though the model can, as we have argued, accommodate such irrationality). Moreover, in our setting, variation in the market's pricing rule is endogenous—there is no analog to exogenous, noise trader-induced sentiment shocks.

Another prominent theme in corporate finance has to do with the consequences of managerial “short-termism.” Work in this area (Narayanan (1985), Stein (1989), Bebchuk and Stole (1993)) has tended to emphasize a relatively static set of inefficiencies that arise when, in a world of asymmetric information, managers are concerned with maximizing current stock prices instead of long-run value. For example, if the stock market focuses exclusively on reported earnings, short term-oriented managers will take various steps to pump up earnings, either via accounting manipulations or through real-side distortions such as cuts in maintenance expenditures, changes in pricing policies designed to bring forward current sales at the expense of future sales, and so forth.

By contrast with this work, our premise is that the market makes inferences about firm value based on multiple variables (i.e., both growth and profitability measures) rather than just earnings. This leads to the potential for a broader set of outcomes than in models of earnings hyping; for example, firms may sometimes invest too much when they are playing the growth strategy. And importantly, multivariate learning generates a more dynamic account of firm behavior, whereby the nature of the distortion varies over time along with the market's approach to valuation.

The one paper in the short-termism/learning genre that is closest to ours is Rajan's (1994) account of bank credit cycles. As we do, Rajan considers the possibility that there will be time variation in how the market makes inferences about managerial ability. In particular, if a bank reports poor performance due to loan losses, it will tend to be evaluated harshly by the market unless many other banks are also suffering loan losses at the same time, in which case the market will blame a systematic shock. This mechanism can lead to cycles in which banks roll over bad loans for several periods in order to hide their losses, and then, when the buildup of bad loans reaches a critical point, they coordinate and all switch to a strategy of loss recognition and credit contraction. This is similar in many ways to our story. Perhaps the most noteworthy difference is that in Rajan's model, the dynamics of learning are driven by a benchmarking externality across banks. That is, there need to be several similar banks, and the market must base its evaluation of bank i in part on the performance of bank j . In our model, on the other hand, each firm is evaluated by the market based solely on its own performance.

Within the large literature on macroeconomic fluctuations, we are probably closest to the work on endogenous business cycles.¹⁷ This research shows that purely intrinsic cycles can arise from a variety of other sources beyond those considered here: (i) the existence of nonlinearities as in Grandmont (1985); (ii) price or wage rigidities that lead to a multiplier–accelerator mechanism as in Goodwin (1951); or (iii) pecuniary externalities combined with credit constraints, as in Aghion, Banerjee, and Piketty (1999). Also related is Shleifer (1986), who shows that endogenous cycles can be driven by each firm’s desire to coordinate the timing of its investments with that of other firms. As with Rajan (1994), Shleifer’s model is thus fundamentally about strategic interaction across firms, whereas ours is about the interplay between a single firm and the stock market.¹⁸

V. Conclusions

The basic message of this paper can be simply stated. A desire on the part of managers to please the stock market can introduce significant excess volatility into real variables such as output and sales, even if the market itself is fully efficient. This conclusion differs from the premise implicit in many recent papers in behavioral finance that the stock market can only exert a destabilizing influence on real activity to the extent that it is subject to nonfundamental sentiment shocks. Our theory makes a number of distinctive empirical predictions, which we hope to see explored in future empirical work.

Appendix A: Endogenizing the Managerial Utility Function

We will couch the argument in terms of the dynamic model; the argument for the two-period static version is just a special case. We begin by making the following assumptions about the timing of actions within each period t :

- (i) First, the firm and the generation- t manager sign an incentive contract. The incentive contract is of the following simple form: The manager is granted shares equivalent to a fractional ownership n of the firm’s stock ($0 \leq n \leq 1$), and these shares may be either all unrestricted or all restricted. If the shares are unrestricted, the manager can sell them at the end of period t if and only if she experiences a liquidity shock—see below. If the shares are restricted, the manager cannot sell them until

¹⁷ The work on endogenous business cycles can be contrasted with the real business cycle (RBC) approach pioneered by Kydland and Prescott (1982), as well as with models based on “sunspots” or “animal spirits” (e.g., Woodford (1990), Howitt and McAfee (1992)). In the RBC framework, fluctuations result from exogenous shocks to, for example, demand or technology. In sunspot models, extrinsic variables also play a key role, namely, that of coordinating agents’ expectations.

¹⁸ Francois and Lloyd-Ellis (2003) embed Shleifer’s (1986) implementation cycles into a full-fledged Schumpeterian framework in which both R&D activities and implementation decisions are endogenous. This means that, as in our model, firms have to divide their efforts between two tasks—in their case, R&D and production.

the end of period $t + 1$. Denote the fraction of shares that are unrestricted by u , where our all-or-none assumption implies that $u \in \{0, 1\}$.

- (ii) Next, the manager chooses an allocation of effort e_t across the growth and margins strategies, as in the basic model in the text.
- (iii) Next, the manager may experience a publicly observable liquidity shock. Specifically, with probability α she learns that she is an early consumer, and can only derive utility from time- t consumption. This implies that at the end of the period, she will sell all the shares she can and will get total utility of $n(\pi_t + uP_t)$. With probability $(1 - \alpha)$, she learns that she is a late consumer and can derive utility from time- $t + 1$ consumption as well. In this case, she is required to hold her shares until time $t + 1$ and she gets total utility of $n(\pi_t + \delta P_{t+1})$, which from her perspective is functionally equivalent to $n(\pi_t + \delta\pi_{t+1})$ since nothing she does has any consequence beyond period $t + 1$.
- (iv) Next, the manager decides whether to exert some further “maintenance” effort to prepare the firm’s assets for period- $t + 1$ production. If this effort is exerted, period- $t + 1$ profits are as given in the text. If not, period- $t + 1$ profits are reduced by a discrete amount X . The maintenance effort costs the manager a monetary equivalent of c .
- (v) Next, the manager’s maintenance effort is observed by the stock market, although it is not directly contractible.
- (vi) Finally, if the manager is an early consumer, she sells any unrestricted shares and consumes the proceeds.

Given these assumptions, the manager understands that if she exerts the maintenance effort in period t , π_{t+1} will be increased by X and P_t will be increased by δX , all else equal. (The latter is because the forward-looking market observes the maintenance decision.) Thus, if she learns that she is an early consumer, she will exert the effort if

$$nu\delta X \geq c. \quad (\text{A1})$$

This expression captures the idea that an early consumer can only be motivated to provide maintenance effort if she is allowed to sell her shares at the end of period t , that is, if these shares are unrestricted. In contrast, a late consumer will exert maintenance effort if

$$n\delta X \geq c. \quad (\text{A2})$$

Intuitively, the restriction parameter u is of no consequence to a late consumer because she always waits until period $t + 1$ to sell her stock.

If X is sufficiently large, the firm will want to elicit maintenance effort from both early and late-consuming managers, in spite of any distortions along the sales and margins dimensions that this may cause. Assume that X is in fact large enough that this is the case. If so, it is clear that the two incentive constraints will bind, so that $u = 1$ and $n = c/\delta X$.

Given this incentive scheme, at the moment when the manager chooses how to allocate effort e_t across the sales and margins strategies, she anticipates

a probability α of receiving a payoff of $n(\pi_t + P_t)$ and a probability $(1 - \alpha)$ of receiving a payoff of $n(\pi_t + \delta\pi_{t+1})$. Thus, her expected payoff is

$$U_t = n(\pi_t + \alpha P_t + (1 - \alpha)\delta\pi_{t+1}). \quad (\text{A3})$$

For the purposes of analyzing her choice of e_t (i.e., of pursuing the growth vs. margins strategies), this induced utility function can be simplified to

$$U_t = \pi_t + \alpha P_t, \quad (\text{A4})$$

since the choice of e_t is independent of the scale n of the compensation package, and, unlike the maintenance decision, does not influence π_{t+1} . This is exactly the managerial utility function assumed in equation (3) for the static model and equation (32) for the dynamic model.

Note, by the way, that stock-based compensation is essential in a setting like this one in which managers have liquidity shocks. If managers were sure to be late consumers, it would be possible to elicit maintenance effort simply by paying the time- t manager as a function of profits eventually realized at time $t + 1$. But when managers are effectively impatient, and have to be rewarded for their efforts before these efforts shows up in profits, the stock price is the only contractible means of doing so.

Appendix B: Proof of Proposition 2, Part (i)

Suppose, hypothetically, that there exists a stationary equilibrium in which managers devote the same interior level of effort e^h to the growth strategy in all periods. In this case, the steady-state value of q_t is given by

$$q^h = \gamma Q / (\gamma + (1 - \gamma)Ae^h). \quad (\text{B1})$$

In order for managers to be willing to allocate an interior level of effort, it must be the case that they do not gain by deviating to either the all-growth or the all-margins strategies. Using arguments similar to those in Section I above, this condition can be shown to imply

$$q^h = \frac{1 + \alpha\delta\beta^{mh}}{1 + \alpha\delta\beta^{sh}}, \quad (\text{B2})$$

where β^{mh} and β^{sh} are the regression coefficients that the market uses to forecast future profits based on observed margins and sales, respectively, given a conjectured effort allocation equal to the hypothesized value of e^h . Intuitively, the larger is β^{mh} relative to β^{sh} , the more weight the market puts on margins relative to sales in the hypothesized equilibrium, and hence the larger must be the size of the market q^h to deter managers from deviating to the all-margins strategy.

The regression coefficients β^{mh} and β^{sh} are both functions of e^h . In particular, we have that

$$\beta^{mh} = ((1 - e^h)v^a + \rho v^m) / ((1 - e^h)^2 v^a + v^m) \quad (\text{B3})$$

and

$$\beta^{sh} = (e^h q^h v^a + \rho v^s) / ((e^h q^h)^2 v^a + v^s). \quad (\text{B4})$$

Using these expressions for β^{mh} and β^{sh} , it can be shown that (B2) implies the restriction that $q^- < q^h < q^+$. In other words, to support an interior level of effort, q^h must lie inside the range in which the static model admits multiple equilibria.

If there is to be a stationary equilibrium, (B1) and (B2) must both be satisfied, which gives us two equations in terms of the two unknowns, q^h and e^h , as well as various other parameters of the model. However, even if there is a solution to these two equations with an interior value of e^h , the resulting stationary equilibrium will not in general be stable.

To see why, suppose that we perturb q_t slightly away from q^h . At this point, there can no longer be a static equilibrium with interior effort. Suppose that the all-growth equilibrium is chosen instead. Then by Assumption 1, we know that the growth equilibrium will be maintained as long as q_t remains above q^- . This in turn implies that q_t will decline over time until either: i) it hits q^- ; or ii) it asymptotes to q^g . Given that $q^- < q^h$, the only way for q_t to converge back to the equilibrium value of q^h is if

$$q^h = q^g. \quad (\text{B5})$$

Thus, for the hypothesized stationary equilibrium to not only exist, but be stable, we require that (B1), (B2), and (B5) all hold. This is an overdetermined system—there are three equations and only two unknowns, q^h and e^h . As a result, a stable stationary equilibrium can only occur on a zero-measure subset of the parameter space. This establishes part (i) of the proposition. As an aside, note that this stability problem does not arise when $\alpha = 0$. In this case $q^- = q^h = 1$, which means that after a perturbation away from the stationary state to the growth equilibrium, we are eventually pushed back to exactly the point where the interior-effort stationary equilibrium is viable. Q.E.D.

REFERENCES

- Aghion, Philippe, Abhijit Banerjee, and Thomas Piketty, 1999, Dualism and macroeconomic volatility, *Quarterly Journal of Economics* 114, 1359–1397.
- Baker, Malcolm, Jeremy C. Stein, and Jeffrey Wurgler, 2003, When does the market matter? Stock prices and the investment of equity-dependent firms, *Quarterly Journal of Economics* 118, 969–1005.
- Baker, Malcolm, and Jeffrey Wurgler, 2004a, A catering theory of dividends, *Journal of Finance* 59, 1125–1165.
- Baker, Malcolm, and Jeffrey Wurgler, 2004b, Appearing and disappearing dividends: The link to catering incentives, *Journal of Financial Economics* 73, 271–288.
- Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor sentiment and the cross section of stock returns, *Journal of Finance* 61, 1645–1680.
- Barberis, Nicholas, and Andrei Shleifer, 2003, Style investing, *Journal of Financial Economics* 68, 161–189.

- Bebchuk, Lucian, and Lars Stole, 1993, Do short-term objectives lead to under- or overinvestment in long-term projects? *Journal of Finance* 48, 719–729.
- Becker, Holly, 2001, Amazon.com: The good, the bad and the ugly, *Lehman Brothers Equity Research*, March 2, p. 3.
- Blanchard, Olivier, Chanyong Rhee, and Lawrence Summers, 1993, The stock market, profit, and investment, *Quarterly Journal of Economics* 108, 115–136.
- Bolton, Patrick, Jose Scheinkman, and Wei Xiong, 2006, Executive compensation and short-termist behavior in speculative markets, *Review of Economic Studies* 73, 577–610.
- Bosworth, Barry, 1975, The stock market and the economy, *Brookings Papers on Economic Activity*, 257–290.
- Chirinko, Robert S., and Huntley Schaller, 2001, Business fixed investment and “bubbles”: The Japanese case, *American Economic Review* 91, 663–680.
- Dewatripont, Matias, Ian Jewitt, and Jean Tirole, 1999, The economics of career concerns, part I: Comparing information structures, *Review of Economic Studies* 66, 183–198.
- Diamond, Douglas, and Philip Dybvig, 1983, Bank runs, deposit insurance and liquidity, *Journal of Political Economy* 91, 401–419.
- Ehrens, Scott, 1999, ROIC is key (not the gross margin), *Bear Stearns Equity Research*, February 12, p. 3.
- Fischer, Stanley, and Robert C. Merton, 1984, Macroeconomics and finance: The role of the stock market, *Carnegie-Rochester Conference Series on Public Policy* 21, 57–108.
- Francois, Patrick, and Huw Lloyd-Ellis, 2003, Animal spirits through creative destruction, *American Economic Review* 93, 530–550.
- Gilchrist, Simon, Charles P. Himmelberg, and Gur Huberman, 2005, Do stock price bubbles influence corporate investment? *Journal of Monetary Economics* 52, 805–827.
- Goodwin, Richard, 1951, The non-linear accelerator and the persistence of business cycles, *Econometrica* 19, 1–17.
- Grandmont, Jean-Michel, 1985, On endogenous competitive business cycles, *Econometrica* 53, 995–1045.
- Holmstrom, Bengt, 1999, Managerial incentive problems: A dynamic perspective, *Review of Economic Studies* 66, 169–182.
- Holmstrom, Bengt, and Paul Milgrom, 1991, Multi-task principal-agent analyses: Incentive contracts, asset ownership and job design, *Journal of Law, Economics and Organization* 7, 24–52.
- Holmstrom, Bengt, and Jean Tirole, 1993, Market liquidity and performance monitoring, *Journal of Political Economy* 101, 678–709.
- Hong, Harrison, and Jeremy C. Stein, 2003, Simple forecasts and paradigm shifts, NBER Working paper 10013.
- Howitt, Peter, and Preston McAfee, 1992, Animal spirits, *American Economic Review* 82, 493–507.
- Keynes, John Maynard, 1936, *The General Theory of Employment, Interest, and Money* (Macmillan: London).
- Kydland, Finn, and Edward Prescott, 1982, Time to build and aggregate fluctuations, *Econometrica* 50, 1345–1370.
- Laffont, Jean-Jacques, and Jean Tirole, 1988, Repeated auctions of incentive contracts, investment, and bidding parity with an application to takeovers, *Rand Journal of Economics* 19, 516–537.
- Malmendier, Ulrike, and Geoffrey Tate, 2005, CEO overconfidence and corporate investment, *Journal of Finance* 60, 2661–2700.
- Miller, Merton H., and Kevin Rock, 1985, Dividend policy under information asymmetry, *Journal of Finance* 40, 1031–1051.
- Morck, Randall, Andrei Shleifer, and Robert W. Vishny, 1990, The stock market and investment: Is the market a sideshow? *Brookings Papers on Economic Activity*, 157–215.
- Mullainathan, Sendhil, 2002, Thinking through categories, Working paper, Harvard University.
- Narayanan, M. P., 1985, Managerial incentives for short-term results, *Journal of Finance* 40, 1469–1484.
- Polk, Christopher, and Paola Sapienza, 2006, The stock market and investment: A test of catering theory, Working paper, Northwestern University.

- Rajan, Raghuram, 1994, Why bank credit policies fluctuate: A theory and some evidence, *Quarterly Journal of Economics* 109, 399–442.
- Shleifer, Andrei, 1986, Implementation cycles, *Journal of Political Economy* 94, 1163–1190.
- Stein, Jeremy C., 1989, Efficient capital markets, inefficient firms: A model of myopic corporate behavior, *Quarterly Journal of Economics* 104, 655–669.
- Stein, Jeremy C., 1996, Rational capital budgeting in an irrational world, *Journal of Business* 69, 429–455.
- Tobin, James, 1969, A general equilibrium approach to monetary theory, *Journal of Money, Credit, and Banking* 1, 15–29.
- Woodford, Michael, 1990, Learning to believe in sunspots, *Econometrica* 58, 277–307.