Growth Versus Margins: Destabilizing Consequences of Giving the Stock Market What it Wants

Citation

Aghion, Philippe, and Jeremy C. Stein. 2008. Growth versus margins: Destabilizing consequences of giving the stock market what it wants. *Journal of Finance* 63, no. 3: 1025-1058.

Published Version

<http://dx.doi.org/10.1111/j.1540-6261.2008.01351.x>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3660730>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

NBER WORKING PAPER SERIES

GROWTH VS. MARGINS:
DESTABILIZING CONSEQUENCES OF
GIVING THE STOCK MARKET WHAT IT WANTS

Philippe Aghion
Jeremy C. Stein

Working Paper 10999
<http://www.nber.org/papers/w10999>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2004

We are grateful to the National Science Foundation for financial support, and to Richard Holden for research assistance. Thanks also to Malcolm Baker, Patrick Bolton, Arvind Krishnamurthy, Tuomo Vuolteenaho, and seminar participants at the NBER for helpful comments. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2004 by Philippe Aghion and Jeremy C. Stein. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Growth vs. Margins: Destabilizing Consequences of Giving the Stock Market What It Wants
Philippe Aghion and Jeremy C. Stein
NBER Working Paper No. 10999
December 2004
JEL No. E32, E44, G30

ABSTRACT

We develop a multi-tasking model in which a firm can devote its efforts either to increasing sales growth, or to improving per-unit profit margins by, e.g., cutting costs. If the firm's manager is concerned with the current stock price, she will tend to favor the growth strategy at those times when the stock market is paying more attention to performance on the growth dimension. Conversely, it can be rational for the stock market to weight observed growth measures more heavily when it is known that the firm is following a growth strategy. This two-way feedback between firms' business strategies and the market's pricing rule can lead to purely intrinsic fluctuations in sales and output, creating excess volatility in these real variables even in the absence of any external source of shocks.

Philippe Aghion
Department of Economics
Harvard University
Littauer Center
Cambridge, MA 02138
and NBER
p_aghion@harvard.edu

Jeremy C. Stein
Department of Economics
Harvard University
Littauer 209
Cambridge, MA 02138
and NBER
jeremy_stein@harvard.edu

1 Introduction

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 four years of single-minded focus on growth, and then just under two 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 multi-tasking problem of the sort envisioned by, e.g., Holmstrom and Milgrom (1991).

In this paper, we study the implications of this multi-tasking 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 various performance measures. In the specific case of Amazon.com, Hong and Stein (2004) 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 changes to being 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.

¹For example, in a February 1999 report lauding Amazon's growth prospects, Scott Ehrens 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..”

²In a March 2001 report on Amazon, Holly Becker 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...”

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 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 the market price is especially sensitive to performance on the growth dimension. Now over time, as it penetrates the market more fully, the firm 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 come too late, it will be too abrupt. Much like Bezos' description of Amazon post-1999, the firm will go to the other extreme of focusing exclusively on cost-cutting, as opposed

³As 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.

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 for too long pay insufficient attention to growth opportunities, 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—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 where 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 cashflows. 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 then 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—i.e., 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 2, 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 3, 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—i.e., when the firm plays the growth strategy for several periods in a row, this reduces the stock of unsatiated cus-

⁴See, e.g., 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 (2004), and Gilchrist, Himmelberg and Huberman (2004).

⁵In 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).

tomers and hence shrinks demand going forward. In Section 4, we flesh out the model’s most distinctive empirical implications. Section 5 discusses related work, and Section 6 concludes.

2 Static Model

2.1 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, it 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 which 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_0 and a margins component m_0 .⁶ Thus:

$$\pi_0 = s_0 + m_0, \quad (1)$$

where s_0 and m_0 depend on both managerial ability, as well as on 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_0 and the remaining $(1 - e)$ of which is devoted to improving margins m_0 . The manager chooses effort so as to maximize her expected utility U , which we take to be equal to a linear combination of current profits and the firm’s stock price:⁷

$$U = E(\pi_0 + \alpha P_0), \quad (2)$$

where the stock price P_0 is equivalent to the market’s discounted expectation of future (i.e., second-period) profits as of the end of the first period. As will become clear shortly, this expectation will depend in part on the market’s updated assessment of managerial ability.

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

$$s_0 = aeq_0 + \varepsilon^s, \quad (3)$$

⁶Naturally, 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.

⁷This sort of managerial utility function appears in a number of previous papers, including, e.g., Miller and Rock (1985), and Stein (1989).

and

$$m_0 = a(1 - e) + \varepsilon^m, \quad (4)$$

where: (i) q_0 denotes the size of the market, which for the time being we take to be exogenously fixed (an assumption we will relax in the next section when we analyze the dynamics of the model); (ii) ε^s is a random variable which captures sales shocks, and is normally distributed with mean zero and variance v^s ; and (iii) ε^m is a random variable which 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, Jewitt and Tirole (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_0 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_0 increases.

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

$$\pi_1 = a + \rho(\varepsilon^s + \varepsilon^m), \quad (5)$$

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.

The timing of actions within the first period is as follows. First, the manager chooses an allocation of effort. 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_0 . Recall that the stock price is simply the discounted expectation of second-period profits, which implies that:

$$P_0 = \delta E(\pi_1 | s_0, m_0, e^*), \quad (6)$$

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.

2.2 First best

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

⁸To 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.

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:

$$\max_e \{E(\pi_0) = aeq_0 + a(1 - e)\}. \quad (7)$$

It then immediately follows that the static first-best is to set $e = 1$ (i.e., to pursue a growth strategy) if $q_0 > 1$ and to set $e = 0$ (i.e., to pursue a margins strategy) if $q_0 < 1$. 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_0 is larger.

2.3 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. Then we 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: $e = e^*$.

2.3.1 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_0 and m_0 , Bayesian updating leads to the following pricing rule for the firm's stock:

$$\begin{aligned} P_0^g &= \delta E(\pi_1 | s_0, m_0, e^* = 1) \\ &= \delta(A + \beta^{sg}(s_0 - Aq_0) + \beta^{mg}m_0), \end{aligned} \quad (8)$$

where, using the fact that the two variables s_0 and m_0 are uncorrelated, we have that:

$$\beta^{sg} = \text{cov}(\pi_1, s_0 | e^* = 1) / \text{var}(s_0 | e^* = 1) = (q_0 v^a + \rho v^s) / (q_0^2 v^a + v^s) \quad (9)$$

and

$$\beta^{mg} = \text{cov}(\pi_1, m_0 | e^* = 1) / \text{var}(m_0 | e^* = 1) = \rho. \quad (10)$$

Similarly, suppose that the market conjectures that the firm is pursuing a margins strategy ($e^* = 0$). Then, based on the observed realizations of s_0 and

m_0 , Bayesian updating leads to the following pricing rule:

$$\begin{aligned} P_0^m &= \delta E(\pi_1 | s_0, m_0, e^* = 0) \\ &= \delta(A + \beta^{sm} s_0 + \beta^{mm} (m_0 - A)), \end{aligned} \quad (11)$$

where

$$\beta^{sm} = \text{cov}(\pi_1, s_0 | e^* = 0) / \text{var}(s_0 | e^* = 0) = \rho \quad (12)$$

and

$$\beta^{mm} = \text{cov}(\pi_1, m_0 | e^* = 0) / \text{var}(m_0 | e^* = 0) = (v^a + \rho v^m) / (v^a + v^m). \quad (13)$$

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_0 is more informative about her ability, and hence about future profits, than is the realization of s_0 . Second, if $\rho q_0 < 1$, (which will always be satisfied in a neighborhood around $q_0 = 1$, 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.

2.3.2 The potential for multiple equilibria

We are now ready to construct the equilibria of the game between the manager and the market. This involves postulating a market conjecture, $e^* \in \{0, 1\}$; then, using the updating rules derived above, computing the expected payoff to the manager under alternative effort allocations; and 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 in fact pursuing the growth strategy ($e = 1$), her expected sales, profits and margins are respectively equal to:

$$E(s_0 | e = 1) = Aq_0 = E(\pi_0 | e = 1); E(m_0 | e = 1) = 0, \quad (14)$$

Therefore, from our analysis in the previous subsection, her expectation of the resulting stock price is equal to:

$$E(P_0^g | e = 1) = \delta(A + \beta^{sg} E(s_0 - Aq_0 | e = 1) + \beta^{mg} E(m_0 | e = 1)) = \delta A. \quad (15)$$

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

$$E(\pi_0 + \alpha P_0^g | e = 1) = Aq_0 + \alpha\delta A. \quad (16)$$

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:

$$E(m_0 | e = 0) = A = E(\pi_0 | e = 0); E(s_0 | e = 0) = 0, \quad (17)$$

and her expectation of the resulting stock price is equal to:

$$\begin{aligned} E(P_0^g | e = 0) &= \delta(A + \beta^{sg} E(s_0 - Aq_0 | e = 0) + \beta^{mg} E(m_0 | e = 0)) \\ &= \delta(A - \beta^{sg} Aq_0 + \beta^{mg} A). \end{aligned} \quad (18)$$

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

$$E(\pi_0 + \alpha P_0^g | e = 0) = A + \alpha\delta(A - \beta^{sg} Aq_0 + \beta^{mg} A). \quad (19)$$

Comparing equations (16) and (19), we see that by deviating away from the market conjecture, the manager generates sales that are below expectations by an average amount Aq_0 , 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:

$$q_0 \geq 1 - \alpha\delta(\beta^{sg} q_0 - \beta^{mg}). \quad (20)$$

Thus for sufficiently large market capacity q_0 , it will always be an equilibrium for the manager to play the growth strategy. Denote by q^- the value of q_0 that satisfies the above constraint with equality. Using the expressions for β^{sg} and β^{mg} derived above, one can analyze how the cut-off point q^- varies with the parameters of the model. In particular, note that q^- solves:

$$q - \alpha\delta\rho = 1 - \alpha\delta \frac{q^2\nu^a + \rho q\nu^s}{q^2\nu^a + \nu^s}. \quad (21)$$

Using this fact, one can show that q^- is a decreasing function of both ability-uncertainty, ν^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 $q_0 \geq 1$.

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:

$$E(\pi_0 + \alpha P_0^m | e = 0) = A + \alpha\delta A. \quad (22)$$

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

$$E(\pi_0 + \alpha P_0^m | e = 1) = Aq_0 + \alpha\delta(A + \beta^{sm}Aq_0 - \beta^{mm}A). \quad (23)$$

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_0 . The manager will choose to conform with the market's expectation and emphasize margins whenever

$$q_0 \leq 1 + \alpha\delta(\beta^{mm} - \beta^{sm}q_0). \quad (24)$$

Denote by q^+ the value of q_0 for which condition (24) holds with equality. Or equivalently call q^+ the solution to:

$$q + \alpha\delta\rho q = 1 + \alpha\delta\frac{\nu^a + \rho\nu^m}{\nu^a + \nu^m}. \quad (25)$$

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

$$q^+ > 1 > q^-, \quad (26)$$

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

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

in the first best case where at least one of ν^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 ν^a are both strictly positive. Then, there exist two cut-off values q^- and q^+ such that: (i) $q^+ > 1 > q^-$; (ii) if $q_0 < q^-$, the unique equilibrium is that the firm plays the margins strategy ($e = 0$); (iii) if $q_0 > q^+$, the unique equilibrium is that the firm plays the growth strategy ($e = 1$); (iv) if $q^- \leq q_0 \leq q^+$, 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 ν^a , but it vanishes as α or ν^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.

3 Dynamics

In the static analysis, we took the market size q_0 to be exogenously fixed. Now we consider the case where 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, e.g., how the volatility of sales growth depends on the various parameters of the model.

3.1 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 which 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 .⁹ 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:

$$q_{t+1} = (1 - \gamma)(1 - A)q_t + \gamma Q_t. \quad (28)$$

In the margins regime, everything is the same, except that per-firm sales at time t are 0. So if the margins regime prevails at time t , we have:

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

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:

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

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.

⁹We 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.

3.2 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:

$$U_t = E_t(\pi_t + \alpha P_t). \quad (31)$$

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

$$P_t = E_t \sum_{k \geq 1} \delta^k \pi_{t+k}, \quad (32)$$

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—i.e., the benefits of her tenure show up in profits even after she is gone. In particular, profits are given by:

$$\pi_t = a_t e_t q_t + a_t (1 - e_t) + \varepsilon_t^s + \varepsilon_t^m + a_{t-1} + \rho(\varepsilon_{t-1}^s + \varepsilon_{t-1}^m), \quad (33)$$

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.¹⁰

$$\pi_t = s_t + m_t, \quad (34)$$

where

$$s_t = a_t e_t q_t + \varepsilon_t^s + \rho \varepsilon_{t-1}^s \quad (35)$$

and

$$m_t = a_t (1 - e_t) + a_{t-1} + \varepsilon_t^m + \rho \varepsilon_{t-1}^m. \quad (36)$$

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

¹⁰In 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.

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 $t + 1$. And finally, at the beginning of period $t + 1$, past 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 (30).

3.3 First best

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.

3.3.1 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. Now, 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 (30) tells us that the dynamics of q_t are given by:

$$q_{t+1} = (1 - \gamma) (1 - A e^f) q_t + \gamma Q. \quad (37)$$

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

$$q^s = \gamma Q / (\gamma + (1 - \gamma) A e^f). \quad (38)$$

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:

$$e^f = \gamma(Q - 1)/((1 - \gamma)A). \quad (39)$$

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 (39) 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 1. And similarly, starting from $q_t < 1$, managers will choose the margins strategy which will make q_t increase towards 1. In contrast, we will show below that when managers do cater to the market, it is in general not possible to sustain a stable steady-state equilibrium of this sort.

If the parameters are such that the right-hand side of (39) does not lie between zero and one, we have a corner solution. In particular, if the right-hand side of (39) 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 (39) is negative, the firm plays the pure margins strategy, and unfilled orders asymptote to Q .

3.3.2 A market with shrinking demand flow

It is straightforward to extend the above reasoning to the case where Q_t declines monotonically. Consider the natural analog to (39) with a time-varying effort e_t^f :

$$e_t^f = \gamma(Q_t - 1)/((1 - \gamma)A). \quad (40)$$

If we start out with a value of Q_t such that the right-hand side of (40) 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 (40) drops below one, at which point we have an interior solution with the level of effort devoted to growth, e_t^f , given by (40). 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 no longer be the case when managers care about stock prices.

3.4 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 cut-off 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 $q_t < q^-$; (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 said differently, whenever there is a choice of equilibria, we select the one played in the prior period.

3.4.1 Within-period analysis

To understand why the same cut-off 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:

$$P_t = \delta E(a_t + \rho(\varepsilon_t^s + \varepsilon_t^m) + \Omega_{t+1}|s_t^a, m_t^a, e_t^*), \quad (41)$$

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} .¹¹ 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_0 is replaced everywhere with q_t .

3.4.2 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:

¹¹In 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.

Assumption A1: 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.¹² 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 non-common-knowledge structure on the market's higher-order beliefs out of equilibrium. In particular, Assumption A1 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, 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.¹³

3.4.3 Deterministic cycles with constant demand flow

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

Proposition 2 *Under Assumption A1, with Q_t constant, and with α and ν^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; (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 $\underline{q} < q^-$.*

¹²Rajan (1994) makes an identical assumption in a model with a similar dynamic structure.

¹³Several readers have suggested another equilibrium-selection criterion: if at date t , there exist multiple equilibria, 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 in which 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 at all 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 in general attainable via a mechanism that relies on managerial announcements, because of the opportunistic nature of these announcements with respect to managers' private information.

Proof: Part (i) of Proposition 2 is proven in the appendix. Here, we just establish part (ii); this part of the proof is briefer and captures most of the important intuition. Let q^g and q^m denote the fixed points of the dynamic equations (28) and (29), i.e., the steady-state values of q_t in the all-growth and all-margins regimes, respectively. We have already noted that:

$$q^g = \gamma Q / (\gamma + (1 - \gamma)A) \quad (42)$$

and:

$$q^m = Q. \quad (43)$$

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

$$q^g < q^- \text{ and } q^m > q^+. \quad (44)$$

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 A1 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 in turn brings us back to the situation where we started—i.e., where q_t belongs to the interval (q^g, q^-) and where 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 (44) 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.

Figure 1 provides an illustration of deterministic cycles with constant demand flow. We choose the following baseline parameter values: $Q = 2$, $\gamma = 0.1$, $A = 0.5$, $\nu^a = \nu^s = \nu^m = 1$, $\rho = 0.1$, 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 (39). 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 case 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—for longer than in the first-best, 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 with $\alpha = 1.5$. However, we now have that $q^- = 0.64$, and $q^+ = 1.59$, i.e., 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.

FIGURE 1 HERE

3.4.4 Non-monotonic adjustment with shrinking demand flow

The other case of interest is the one where the demand flow declines monotonically over time. As seen above, in the first-best 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 non-monotonic.

The mechanism is very similar to that which 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:

$$Q_t = Q_0 \beta^t, \quad (45)$$

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:

$$q_{t+1} = (1 - \gamma)(1 - A)q_t + \gamma Q_0 \beta^t. \quad (46)$$

According to (46), q_t will eventually start declining. Moreover, by Assumption A1, 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, 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:

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

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 regime. But what is interesting relative to the first-best case is that this convergence need not be monotonic. Indeed, it can in principle 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 non-monotonic convergence. We set $Q_0 = 20$, $\beta = 0.9$, $\gamma = 0.1$, $A = 0.5$, $\nu^a = \nu^s = \nu^m = 1$, $\rho = 0.1$, 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 (40). 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. Therefore the firm 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 for 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, goes back to growth, all before settling permanently into the margins strategy.

FIGURE 2 HERE

3.5 Behavioral variations

3.5.1 Investors use simple models

So far we have taken stock-market investors to be fully rational—i.e., 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 either: i) current profits; ii) current sales; or iii) 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 non-zero. 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_1 in the growth regime is s_0 , and the best predictor of π_1 in the margins regime is m_0 . Moreover, since s_0 and m_0 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 2, we can calculate cut-off values of q_0 that define the regions over which the growth and margins equilibria can exist. Denoting these revised cutoffs by q_b^- and that q_b^+ , we have that q_b^- solves:

$$q = 1 - \alpha\delta \frac{q^2\nu^a + \rho q\nu^s}{q^2\nu^a + \nu^s}. \quad (48)$$

This equation is almost identical to its counterpart (21) 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:

$$q_b^- < q^-. \quad (49)$$

Similarly, q_b^+ satisfies:

$$q = 1 + \alpha\delta \frac{\nu^a + \rho\nu^m}{\nu^a + \nu^m}. \quad (50)$$

This equation also looks like its counterpart (25) in the rational case, except for the term $(\alpha\delta\rho q)$ which has disappeared from the left hand side. This in turn yields:

$$q_b^+ > q^+. \quad (51)$$

¹⁴This 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 (2004) for an extensive discussion, and references to the relevant psychology literature.

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.

When we move to the dynamic setting, everything continues to work exactly as before, except that we now use the new cut-off 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 $Q = 2$, $\gamma = 0.1$, $A = 0.5$, $\nu^a = \nu^s = \nu^m = 1$, $\rho = 0.1$, $\delta = 1$, 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.

FIGURE 3 HERE

3.5.2 Investors treat stocks in the same category similarly

Another potentially interesting extension of our model—which we do not pursue formally here—flows 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 a 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.

4 Empirical Implications

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

4.1 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—i.e., the *sensitivity* of stock prices to observable variables such as sales and profit margins. Thus our central implication is 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 use 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 over time to changes in the nature of firm performance. In particular, one could ask whether, e.g., investment is on average higher for firms in industry i at those times when the industry’s growth beta is larger. And with the appropriate data on managerial incentives, one could further check to see if this relationship is more pronounced in the subset of firms where managers have a lot of stock or options that can be sold in the near future.

We are unaware of any empirical work that explores these questions. However, we should mention a pair of recent papers by Baker and Wurgler (2004a, 2004b) 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 non-dividend-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 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.¹⁵

4.2 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.

¹⁵One distinction between our work and Baker and Wurgler (2004a, 2004b) 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.

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 (2004c). 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.

5 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. Like 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.¹⁶ This research shows that purely intrinsic cycles can arise from a variety of other sources beyond those considered here: (i) the existence of non-linearities as in Grandmont (1985); (ii) price or wage rigidities which 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.¹⁷

6 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

¹⁶The 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, e.g., demand or technology. In sunspot models, extrinsic variables also play a key role, namely that of coordinating agents’ expectations.

¹⁷Francois 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.

into real variables such as output and sales, even if the market itself is fully efficient. This conclusion differs from the premise implicit in much recent research in behavioral finance—that the stock market can only exert a destabilizing influence on real activity to the extent that it is subject to non-fundamental sentiment shocks. Our theory makes a number of distinctive empirical predictions, which we hope to explore in future empirical work.

7 Appendix

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:

$$q^h = \gamma Q / (\gamma + (1 - \gamma)Ae^h), \quad (52)$$

In order for managers to be willing to allocate an interior level of effort, it must be that they do not gain by deviating to either the all-growth or the all-margins strategies. Using arguments similar to those in Section 2 above, this condition can be shown to imply:

$$q^h = \frac{1 + \alpha\delta\beta^{mh}}{1 + \alpha\delta\beta^{sh}} \quad (53)$$

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:

$$\beta^{mh} = ((1 - e^h)v^a + \rho v^m) / ((1 - e^h)^2 v^a + v^m). \quad (54)$$

and:

$$\beta^{sh} = (e^h q^h v^a + \rho v^s) / ((e^h q^h)^2 v^a + v^s). \quad (55)$$

Using these expressions for β^{mh} and β^{sh} , it can be shown that (53) 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 where the static model admits multiple equilibria.

If there is to be a stationary equilibrium, (52) and (53) 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 A1, 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:

$$q^h = q^g. \tag{56}$$

Thus for the hypothesized stationary equilibrium to not only exist, but be stable, we require that (52), (53), and (56) 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.

References

- [1] Aghion, Philippe, Abhijit Banerjee, and Thomas Piketty, 1999, "Dualism and Macroeconomic Volatility," *Quarterly Journal of Economics* 114, 1359-1397.
- [2] 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.
- [3] Baker, Malcolm, and Jeffrey Wurgler, 2004a, "A Catering Theory of Dividends," *Journal of Finance* 59, 1125-1165.
- [4] Baker, Malcolm, and Jeffrey Wurgler, 2004b, "Appearing and Disappearing Dividends: The Link to Catering Incentives," *Journal of Financial Economics* (forthcoming).
- [5] Baker, Malcolm, and Jeffrey Wurgler, 2004c, "Investor Sentiment and the Cross Section of Stock Returns," NBER working paper 10449.
- [6] Barberis, Nicholas and Andrei Shleifer, 2003, "Style Investing," *Journal of Financial Economics* 68, 161-189.
- [7] 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.
- [8] Becker, Holly, 2001, "Amazon.com: The Good, the Bad and the Ugly," *Lehman Brothers Equity Research*, p. 3 (March 2).
- [9] Blanchard, Olivier, Chanyong Rhee, and Lawrence Summers, 1990, "The Stock Market, Profit, and Investment," *Quarterly Journal of Economics* 105, 115-136.
- [10] Bosworth, Barry, 1975, "The Stock Market and the Economy," *Brookings Papers on Economic Activity*, 257-290.
- [11] Chirinko, Robert S., and Huntley Schaller, 2001, "Business Fixed Investment and 'Bubbles': The Japanese Case," *American Economic Review* 91, 663-680.
- [12] Day, Richard, 1982, "Irregular Growth Cycles," *American Economic Review* 72, 406-414.
- [13] 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.
- [14] Ehrens, Scott, 1999, "ROIC is Key (Not the Gross Margin)," *Bear Stearns Equity Research*, p. 3 (February 12).

- [15] 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.
- [16] Francois, Patrick, and Huw Lloyd-Ellis, 2003, "Animal Spirits Through Creative Destruction," *American Economic Review* 93, 530-550.
- [17] Gilchrist Simon, Charles P. Himmelberg, and Gur Huberman, 2004, "Do Stock Price Bubbles Influence Corporate Investment?," Federal Reserve Bank of New York staff report 177.
- [18] Goodwin, Richard, 1951, "The Non-Linear Accelerator and the Persistence of Business Cycles," *Econometrica* 19, 1-17.
- [19] Grandmont, Jean-Michel, 1985, "On Endogenous Competitive Business Cycles," *Econometrica* 53, 995-1045.
- [20] Holmstrom, Bengt, 1999, "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies* 66, 169-182.
- [21] 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.
- [22] Holmstrom, Bengt, and Jean Tirole, 1993, "Market Liquidity and Performance Monitoring," *Journal of Political Economy* 101, 678-709.
- [23] Hong, Harrison, and Jeremy C. Stein, 2004, "Simple Forecasts and Paradigm Shifts," Harvard University working paper.
- [24] Howitt, Peter, and Preston McAfee, 1992, "Animal Spirits," *American Economic Review* 82, 493-507.
- [25] Keynes, John Maynard, 1936, *The General Theory of Employment, Interest, and Money*, (London: Macmillan).
- [26] Kydland, Finn, and Edward Prescott, 1982, "Time to Build and Aggregate Fluctuations," *Econometrica* 50, 1345-1370.
- [27] Miller, Merton H., and Kevin Rock, 1985, "Dividend Policy Under Information Asymmetry," *Journal of Finance* 40, 1031-1051.
- [28] Morck, Randall, Robert W. Vishny, and Andrei Shleifer, 1990, "The Stock Market and Investment: Is the Market a Sideshow?," *Brookings Papers on Economic Activity*, 157-215.
- [29] Mullainathan, Sendhil, 2002, "Thinking Through Categories," Harvard working paper.
- [30] Narayanan, M.P., 1985, "Managerial Incentives for Short-Term Results," *Journal of Finance* 40, 1469-1484.

- [31] Polk, Christopher, and Paola Sapienza, 2004, "The Real Effects of Investor Sentiment," Northwestern University working paper.
- [32] Rajan, Raghuram, 1994, "Why Bank Credit Policies Fluctuate: A Theory and Some Evidence," *Quarterly Journal of Economics* 109, 399-442.
- [33] Shleifer, Andrei, 1986, "Implementation Cycles," *Journal of Political Economy* 94, 1163-1190.
- [34] Stein, Jeremy C., 1989, "Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior," *Quarterly Journal of Economics*, 104, 655-669.
- [35] Stein, Jeremy C., 1996, "Rational Capital Budgeting in an Irrational World," *Journal of Business* 69, 429-455.
- [36] Tobin, James, 1969, "A General Equilibrium Approach to Monetary Theory," *Journal of Money, Credit, and Banking* 1, 15-29.W
- [37] Woodford, Michael, 1990, "Learning to Believe in Sunspots," *Econometrica* 58, 277-307.

Figure 1: Deterministic Cycles With Constant Demand Flow

The parameter values are chosen as follows: $Q_0 = 2, \gamma = 0.1, A = 0.5, v^a = v^s = v^m = 1, \rho = 0.1$ and $\delta = 1$. We then plot the evolution of q_t for three cases (i) $\alpha = 0$, (ii) $\alpha = 1.0$, (iii) $\alpha = 1.5$.

Figure 2: Non-Monotonic 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) $\alpha = 0$, (ii) $\alpha = 1.0$.

Figure 3: The Effect of Bounded Rationality With Constant Demand Flow

The parameter values are chosen as follows: $Q_0 = 2, \gamma = 0.1, A = 0.5, v^a = v^s = v^m = 1, \rho = 0.1$ and $\delta = 1$. We then plot the evolution of q_t for two cases (i) the case where the investors are fully rational, and (ii) the case where the investors are boundedly rational.