

Entrepreneurial Learning, the IPO Decision, and the Post-IPO Drop in Firm Profitability

Ľuboš Pástor

Lucian Taylor

Pietro Veronesi*

September 24, 2007

Abstract

We develop a dynamic model of the optimal IPO decision in the presence of learning about the average profitability of a private firm. In the IPO decision, the entrepreneur trades off diversification benefits of going public against benefits of private control. The model predicts that firm profitability should decline after the IPO, on average, and that this decline should be larger for firms with more volatile profitability and firms with less uncertain average profitability. These predictions are supported empirically in a sample of 7,183 IPOs in the U.S. between 1975 and 2004.

*All authors are at the Graduate School of Business, University of Chicago. Pástor and Veronesi are also affiliated with the CEPR and NBER. We thank Ray Ball, Markus Brunnermeier (discussant), Doug Diamond, Michael Halling (discussant), John Heaton, Gene Kandel, Peter Kondor, Morten Sorensen, Matt Spiegel, Ivo Welch (discussant), Frank Yu, Mike Weisbach, an anonymous referee, and the audiences at the NBER Entrepreneurship Meeting (Spring 2007), the NBER Asset Pricing Meeting (Spring 2007), the Vienna Symposia in Asset Management, New York University, University of Chicago, University of Illinois at Urbana-Champaign, and University of North Carolina for helpful comments.

1. Introduction

The decision to go public is one of the most important decisions made by private firms. This decision can have various motives, such as to diversify the entrepreneur's holdings, to raise capital for investment, to exploit favorable market conditions, to facilitate acquisitions, to find the firm's market value, to make the firm's shares more liquid, and to make the firm more visible¹. One complicating factor in the IPO decision is that the private firm's future cash flow is highly uncertain, which makes the firm's valuation difficult. We examine the effect of this uncertainty on the IPO decision and on firm profitability around the IPO.

We develop a dynamic model of the optimal IPO decision in the presence of learning about average profitability. The profitability of a private firm follows a process with an unknown mean. Agents learn about this mean by observing realized profits. There are two types of risk-averse agents: investors, who are well diversified, and an entrepreneur, whose entire wealth is tied up in the private firm. The entrepreneur suffers from under-diversification but enjoys benefits of private control. If he takes his firm public, he forfeits these benefits but achieves better diversification by investing the IPO proceeds in stocks and bonds.

The diversification benefits of an IPO are twofold. The first benefit is risk reduction, since a portfolio of stocks and bonds is less risky than concentrated private firm holdings. The second benefit is consumption smoothing. When the private firm's expected profitability rises, the entrepreneur expects higher consumption in the future. He wants to increase his consumption immediately but he cannot borrow against expected future profitability. If expected profitability rises high enough, the entrepreneur's consumption path under private ownership becomes so unattractively steep that he prefers to cash out in an IPO and smooth his consumption by trading stocks and bonds. In short, going public allows the entrepreneur to smooth his consumption not only across states of nature but also across time.

The model produces a cutoff rule whereby going public is optimal when the firm's expected future profitability is sufficiently high. An IPO is optimal when the market value of the firm (value to investors) exceeds the private value of the firm (value to the entrepreneur). Since the entrepreneur's consumption under private ownership is derived from the firm's assets in place rather than from future growth opportunities, the firm's private value is less sensitive to expected future profitability than the firm's market value is. When expected profitability rises, the market value rises faster than the private value, and when expected profitability rises high enough, it becomes optimal for the firm to be owned publicly (by investors) rather than privately (by the entrepreneur).

¹Ritter and Welch's (2002) survey discusses the various motives for going public. For a recent study of the broader security issuance decision, see Henderson, Jegadeesh, and Weisbach (2006).

The model predicts that firm profitability should drop after the IPO, on average, and that this drop should be larger for firms with more volatile profitability and firms with less uncertain average profitability. These novel predictions follow from Bayesian learning and the endogeneity of the IPO. For an IPO to take place, the agents' expected profitability must go up before the IPO, as explained in the previous paragraph. According to Bayes' rule, agents revise their expectations upward only if they observe realized profitability that is higher than expected. As a result, realized profitability exceeds expected profitability at the time of the IPO, and hence profitability is expected to drop after the IPO.

The implications for volatility and uncertainty also follow from Bayesian updating. Agents revise their expectations by less if their prior uncertainty is lower (because prior beliefs are stronger) and if signal volatility is higher (because the signals are less precise). In both cases, realized profitability must rise more sharply in order to pull expected profitability above the IPO cutoff. As a result, the expected post-IPO drop in profitability is larger when volatility is higher and when uncertainty is lower.

To analyze the quantitative predictions of our model, we calibrate the model and compute the expected post-IPO drop in profitability, using a closed-form solution for this expectation. We find that the model delivers empirically plausible dynamics of profitability, both before and after the IPO, including realistic magnitudes for the post-IPO drop.

We test the model's predictions empirically in a sample of 7,183 IPOs in the U.S. between 1975 and 2004. Our evidence supports the model. Firm profitability, measured as return on equity (ROE), declines significantly after the IPO. The average decline in quarterly ROE is 2.7% after one year and 4.2% after three years. A post-IPO drop in profitability has already been reported by Degeorge and Zeckhauser (1993), Jain and Kini (1994), Mikkelson, Partch, and Shah (1997), and Pagano, Panetta, and Zingales (1998) but our sample is much larger.² More important, we also find that the post-IPO drop is larger for stocks with more volatile profitability and firms with less uncertain average profitability. These findings, which do not seem to appear in the literature, support our learning model. They also help distinguish our model from alternative hypotheses that do not involve Bayesian learning.

While the volatility of profitability can be estimated directly from realized profits, uncertainty about average profitability is more difficult to measure. The common proxies for uncertainty also proxy for volatility. To separate uncertainty from volatility, we estimate the stock price reaction to earnings announcements, which should be stronger for firms with

²Degeorge and Zeckhauser (1993) analyze 62 reverse leveraged buyouts (LBOs) in 1983–1987, Jain and Kini (1994) study 682 IPOs in 1976–1988, Mikkelson, Partch, and Shah (1997) examine 283 IPOs in 1980–1983, and Pagano, Panetta, and Zingales (1998) investigate 69 Italian IPOs in 1982–1992.

higher uncertainty and lower volatility. We find that firms with weaker price reactions tend to experience larger post-IPO drops in ROE, as predicted by the model.

The model also predicts that firm profitability increases before the IPO. We do not test this prediction due to the lack of reliable pre-IPO data, but supporting evidence is provided by Degeorge and Zeckhauser (1993) who study 62 reverse LBOs that went public between 1983 and 1987. They find that profitability increases sharply before LBOs return to public ownership and decreases thereafter, consistent with our model. Chemmanur, He, and Nandy (2006) estimate total factor productivity for a large sample of public and private U.S. manufacturing firms. They find that productivity increases steadily in the five years prior to the IPO and decreases steadily thereafter. To the extent that productivity and profitability are comparable, this evidence also supports our model.

There is no asymmetric information in our model. Other models assume, quite plausibly, that the entrepreneur has private information about his firm (e.g., Stein, 1989, Chemmanur, 1993, Chemmanur and Fulghieri, 1999, Maksimovic and Pichler, 2001). Although some information about IPOs is surely asymmetric, it seems useful to know which facts about IPOs can be explained in the more parsimonious world of symmetric information. We show that learning under symmetric information can get us far in understanding IPO profitability. Asymmetric information may well explain some fraction of the observed post-IPO drops in profitability, but learning also seems to play its own distinct role. It is not clear how asymmetric information without learning would generate larger drops in profitability for firms with more volatile profits and firms with less uncertain average profits.

Another possible explanation for the observed profitability pattern is earnings management. Teoh, Welch, and Wong (1998) argue that firms opportunistically inflate their earnings through discretionary accruals shortly before going public. The earnings management hypothesis and our learning hypothesis are not mutually exclusive, but the predictions regarding volatility and uncertainty seem unique to the learning hypothesis. In fact, we explain later that the earnings management hypothesis appears to make the opposite prediction regarding volatility, contrary to the empirical evidence. We also find that our empirical results are robust to controlling for earnings management via discretionary current accruals.

The key motive for an IPO in our model is diversification. This motive is empirically important according to Bodnaruk et al (2006), who find that firms held by less diversified shareholders are more likely to go public. In Benninga, Helmantel, and Sarig (2005), the IPO decision is also driven by the tradeoff between diversification benefits and private benefits. There are important differences between their paper and ours. First, the models are

different. In their model, there is no learning, and agents’ preferences are not modeled explicitly. In our model, learning is crucial, and preferences are modeled in a way that clarifies the diversification benefits of an IPO (risk reduction and consumption smoothing). Second, Benninga et al do not examine IPO profitability, which is the focus of our analysis. Finally, their contribution is theoretical whereas ours is both theoretical and empirical. Diversification benefits of an IPO also feature in Pagano (1993) and Chemmanur and Fulghieri (1999). In the latter model, these benefits are weighed against the lower information production costs associated with private ownership. Information is asymmetric and learning is costly, unlike in our model. Chemmanur and Fulghieri analyze how information production affects IPO timing, whereas we focus on how learning affects IPO profitability.

This paper is also related to the theory of “rational IPO waves” of Pástor and Veronesi (2005). In their model, the entrepreneur observes time-varying market conditions before deciding when to go public. IPO waves arise because many entrepreneurs find it optimal to go public after market conditions improve.³ Unlike in that model, we hold market conditions constant, for simplicity, and focus instead on learning about the private firm itself. In their model, the IPO proceeds are invested in the firm, whereas in our model, they are invested in stocks and bonds for diversification reasons. While they focus on optimal IPO timing, we focus on the dynamics of profitability around the IPO.

Clementi (2002) builds a dynamic model in which firms go public after positive and persistent productivity shocks. Like our model, Clementi’s model implies that firm productivity declines after the IPO, but the mechanism is very different – in Clementi’s model, it is decreasing returns to scale, whereas in our model, it is learning (we have constant returns to scale). Since Clementi’s model features no learning, it does not make the same predictions as our model regarding volatility and uncertainty. On the other hand, Clementi’s model makes some predictions that our model does not make. The main motive for an IPO in Clementi’s model, and also in the model of Pástor and Veronesi (2005), is to raise capital for investment. Both of those models predict an increase in investment after the IPO, which we indeed observe in the data (e.g., Jain and Kini, 1994, and Chemmanur et al, 2006). In contrast, our model makes no interesting predictions for investment because the IPO proceeds are not invested in the firm’s production but rather in stocks and bonds.⁴

Pagano, Panetta, and Zingales (1998) find that the main factor affecting the probability of an IPO is the industry’s market-to-book ratio: An IPO is more likely when the market-

³Consistent with this argument, CFOs identify overall stock market conditions as “the single most important determinant of timing” of an IPO in Brau and Fawcett’s (2006) survey.

⁴When we extend the model to allow some of the IPO proceeds to increase the firm’s investment, we find the same predictions for the dynamics of profitability around the IPO. See Section 4.2.

to-book ratio is high. Pagano et al attribute this finding to the entrepreneurial ability to exploit mispricing, mainly because they also find that firm profits tend to be high before the IPO and lower after the IPO. Mispricing is a possible interpretation, but all three findings are also consistent with our rational model. In the model, all agents find it optimal for firms to go public when the expected profitability is high, and high expected profitability implies a high market-to-book. All of our agents also expect profitability to rise before the IPO and drop after the IPO. Our model is also consistent with the findings of Pagano et al that firms use the IPO proceeds mostly to rebalance their accounts rather than to finance investments, and that IPOs tend to be followed by an increased turnover in control.

The paper is organized as follows. Section 2 presents a toy model that illustrates our basic learning mechanism. Section 3 develops the full model. Section 4 analyzes the model-implied dynamics of profitability, with a focus on the expected post-IPO drop in profitability. Section 5 tests the model’s implications empirically. Section 6 concludes.

2. A Toy Model

In this section, we present a simple model that demonstrates the effect of learning on the behavior of profitability after an IPO. There are two periods, 0 and 1, in which an entrepreneur decides whether to take his private firm public. This decision is made based on a cutoff rule: an IPO takes place if the firm’s expected profitability exceeds a given cutoff. (This type of rule is shown to be optimal in the full model in Section 3.) Let $\underline{\rho}$ denote the cutoff, which is known, and $\bar{\rho}$ denote the firm’s average profitability, which is unknown.

At time 0, the entrepreneur’s prior beliefs about $\bar{\rho}$ are given by the normal distribution,

$$\bar{\rho} \sim N(\hat{\rho}_0, \hat{\sigma}_0^2). \tag{1}$$

At time 1, the entrepreneur observes a signal about average profitability $\bar{\rho}$, namely realized profitability ρ , whose distribution conditional on $\bar{\rho}$ is given by

$$\rho \sim N(\bar{\rho}, \sigma_\rho^2). \tag{2}$$

Result 1. Firm profitability is expected to fall after an IPO at time 1.

To prove this result, we first compute the entrepreneur’s posterior beliefs after observing the signal. Using Bayes’ rule, the posterior distribution of $\bar{\rho}$ is given by

$$\bar{\rho} \mid \rho \sim N(\hat{\rho}, \hat{\sigma}^2), \tag{3}$$

where

$$\hat{\rho} = w_0 \hat{\rho}_0 + (1 - w_0) \rho \quad (4)$$

$$w_0 = \frac{1/\hat{\sigma}_0^2}{1/\hat{\sigma}_0^2 + 1/\sigma_\rho^2}. \quad (5)$$

An IPO takes place at time 1 if expected profitability exceeds the cutoff $\underline{\rho}$:

$$\hat{\rho} \geq \underline{\rho}. \quad (6)$$

Since the IPO takes place at time 1, there is no IPO at time 0, so that

$$\hat{\rho}_0 < \underline{\rho}. \quad (7)$$

Combining equations (6) and (7), we have $\hat{\rho} > \hat{\rho}_0$. It then follows from equation (4) that

$$\rho > \hat{\rho}. \quad (8)$$

In words, for an IPO to take place at time 1, realized profitability ρ must exceed expected profitability $\hat{\rho}$. As a result, the post-IPO profitability is expected to be lower than ρ . At time 0, the expected post-IPO drop in profitability is $E_0(\rho - \hat{\rho} \mid \text{IPO at time 1}) > 0$. A simple and intuitive graphical representation of Result 1 is available in Figure 1.

To simplify the algebraic exposition, add the assumption that $\hat{\rho}_0 = 0$.

Result 2. The post-IPO drop in profitability is expected to be large when the volatility of profitability (σ_ρ) is high and when prior uncertainty about average profitability ($\hat{\sigma}_0$) is low.

To prove this result, rewrite equation (4) as $\rho - \hat{\rho} = w_0 (\rho - \hat{\rho}_0)$. The assumption $\hat{\rho}_0 = 0$ implies $\rho > 0$, so the time-0 expectation of the post-IPO percentage drop in profitability is

$$E_0 \left(\frac{\rho - \hat{\rho}}{\rho} \mid \text{IPO at time 1} \right) = w_0. \quad (9)$$

From equation (5), w_0 increases with σ_ρ and decreases with $\hat{\sigma}_0$. As a result, the expected percentage drop in profitability after the IPO is high when profitability is highly volatile and when there is low uncertainty about average profitability.

The intuition behind both results is simple. For an IPO to take place at time 1, expected profitability must go up between times 0 and 1, so realized profitability at time 1 must exceed expected profitability to “pull it up” via Bayesian updating. Since realized profitability exceeds expected profitability at the IPO, profitability is expected to fall after the IPO (Result 1). If volatility is higher, realized profitability is a less precise signal, so it must

rise by more to pull expected profitability above the IPO cutoff. Similarly, if uncertainty is lower, realized profitability must rise by more to overcome stronger prior beliefs. In both cases, the gap between realized and expected profitability widens, so the post-IPO drop in profitability is larger (Result 2). This intuition applies not only to the percentage drop but also to the absolute drop in profitability. Note that our arguments rely only on Bayesian updating (equations (3) to (5)), the endogeneity of the IPO decision (equation (6)), and the endogeneity of the private firm’s existence before the IPO (equation (7)). We do not assume mean reversion in profitability; profitability in each period is i.i.d., following equation (2).

3. The Full Model

In this section, we develop a model that is richer than our toy model, with more realistic dynamics for profitability and additional assumptions about agent preferences and investment opportunities. This “full” model serves four purposes. First and foremost, it shows that a version of the IPO cutoff rule in equation (6) arises endogenously. Second, it can be calibrated with realistic parameters to deliver empirically plausible dynamics of profitability around the IPO. Third, it allows us to examine the implications of optimal IPO timing. Finally, it motivates the proxy for uncertainty that we use in the empirical analysis.

There are two types of agents, investors and an entrepreneur, and two publicly-traded assets, risky “stocks” and risk-free “bonds”. At time 0, investors are endowed with a large amount of stocks and bonds. The entrepreneur is endowed with a private firm implementing a patent-protected technology. All of the entrepreneur’s wealth is invested in the firm. The entrepreneur has an option to take the firm public at a given time τ , $0 < \tau < T$.⁵ The IPO decision is irreversible. If the entrepreneur chooses to go public, he sells the firm to investors for its fair market value and invests the proceeds in stocks and bonds.

The private firm uses capital B_t to produce earnings at the rate Y_t . The firm’s profitability $\rho_t = Y_t/B_t$ follows the mean-reverting process

$$d\rho_t = \phi(\bar{\rho} - \rho_t) dt + \sigma_{\rho,1}dX_{1,t} + \sigma_{\rho,2}dX_{2,t}, \quad 0 \leq t \leq T, \quad (10)$$

where $\bar{\rho}$ denotes average profitability, ϕ denotes the speed of mean reversion, and $X_{1,t}$ and $X_{2,t}$ are uncorrelated Brownian motions that capture systematic ($X_{1,t}$) and firm-specific ($X_{2,t}$) shocks to firm profitability.⁶ The firm reinvests all of its earnings, so its book value B_t follows the process $dB_t = Y_t dt = \rho_t B_t dt$. The patent expires at time T , at which point the firm’s market value equals the book value, $M_T = B_T$.⁷ Note that since ρ_t follows a random

⁵The model’s key predictions are the same when τ is chosen optimally, as we show in Section 4.2.

⁶Empirically, firm profitability is mean-reverting, e.g., Beaver (1970) and Fama and French (2000).

⁷See Pástor and Veronesi (2003) for a more detailed justification of the terminal value assumption.

process (equation 10), B_t also evolves randomly, so B_T is unknown before time T .

Both the entrepreneur and investors are fully rational utility-maximizing agents. Investor preferences are characterized by a pricing kernel π_t , which follows the stochastic process⁸

$$\frac{d\pi_t}{\pi_t} = -r dt - \sigma_{\pi,1} dX_{1,t}, \quad (11)$$

where r is the risk-free rate and $dX_{1,t}$ is perfectly correlated with the return on stocks. The entrepreneur's preferences at time t are given by

$$\max E_t \left[\int_t^T e^{-\beta(u-t)} \frac{c_u^{1-\gamma}}{1-\gamma} du + \eta e^{-\beta(T-t)} \frac{W_T^{1-\gamma}}{1-\gamma} \right] \quad (12)$$

where c_u denotes consumption, W_T is the entrepreneur's terminal wealth, β is the intertemporal discount, and the local curvature $\gamma > 1$. The entrepreneur retires at time T .

As long as the entrepreneur owns the private firm, he consumes benefits of private control. These benefits include any costs saved by a firm that is not publicly traded (e.g., the costs of separating ownership from control, reporting costs, administrative costs, auditing costs, etc.) as well as benefits commonly referred to as private benefits of control (e.g., Dyck and Zingales, 2004). We distinguish benefits of private control from private benefits of control because the latter benefits can be consumed not only by entrepreneurs but also by managers of publicly traded firms. The benefits of private control disappear immediately after the IPO. For simplicity, we assume that the consumption flow from benefits of private control is proportional to the size of the firm as measured by assets in place,

$$c_t = \alpha B_t, \quad (13)$$

and that the entrepreneur consumes nothing else while managing the private firm. This simplifying assumption is motivated by the fact that entrepreneurs tend to have limited control over their consumption paths due to various constraints, such as borrowing constraints.⁹

There is no asymmetric information. Average profitability \bar{p} in equation (10) is unknown to all agents, investors and entrepreneurs alike. All other parameters are known. Agent beliefs about \bar{p} at time $t = 0$ are represented by the normal prior distribution,

$$\bar{p} \sim N(\hat{\rho}_0, \hat{\sigma}_0^2). \quad (14)$$

⁸A pricing kernel, also called the stochastic discount factor, serves to discount future cash flow in computing its present value. The present value at time s of a future cash flow C_t at time t is $E((\pi_t/\pi_s)C_t)$.

⁹The assumption (13) helps us solve the model analytically. In Section 4.2., we relax this assumption and let the entrepreneur solve the optimal consumption-investment problem before the IPO. We find that the optimal consumption is close to that in (13), and that the model's key predictions remain unchanged.

All agents observe realized profitability ρ_t as well as π_t and they update their beliefs about $\bar{\rho}$ dynamically following Bayes' rule. The agents' posterior beliefs at time t are given by

$$\bar{\rho} \sim N(\hat{\rho}_t, \hat{\sigma}_t^2). \quad (15)$$

The posterior mean $\hat{\rho}_t$ and variance $\hat{\sigma}_t^2$ are given in equation (27) in the Appendix.

3.1. Solving the Model

In this section, we outline the solution of the model and highlight the intuition behind the results. The technical details of the solution are provided in the Appendix.

At time τ , the entrepreneur must decide whether to take his private firm public. He makes this decision by comparing two utility values:

1. The utility $V(M_\tau, \tau)$ resulting from taking the firm public at time τ and investing the IPO proceeds in publicly-traded stocks and bonds until time T
2. The utility $V^O(B_\tau, \tau)$ resulting from keeping the firm private between times τ and T

If the entrepreneur chooses to go public, he can sell any fraction of the firm in the IPO. We show in the Appendix that it is optimal to sell the whole firm, essentially because the benefits of private control disappear after the IPO.¹⁰ The IPO proceeds are thus equal to the firm's market value, M_τ . The outside investors value the firm as the present value of B_T , so the market value is given by $M_\tau = E_\tau[\pi_T B_T] / \pi_\tau$. The closed-form solution for M_τ is in equation (28) in the Appendix. We compute the utility value of M_τ under the assumption that this wealth is optimally invested in stocks and bonds. We do this by transforming the dynamic consumption-investment optimization problem into a static problem (Cox and Huang, 1989). The value function $V(M_\tau, \tau)$ is given in equation (31) in the Appendix.

If the entrepreneur chooses not to go public, he continues consuming benefits of private control, and his terminal wealth is equal to B_T . Using equations (12) and (13), we solve for the utility function $V^O(B_\tau, \tau)$, and report it in equation (32) in the Appendix.

The entrepreneur will take his firm public at time τ if and only if

$$V(M_\tau, \tau) \geq V^O(B_\tau, \tau). \quad (16)$$

Let $P_\tau = V^{-1}(V^O(B_\tau, \tau), \tau)$ define the firm's "private value" at time τ . The entrepreneur is indifferent between owning the private firm and having P_τ dollars optimally invested in

¹⁰In reality, the entrepreneur often retains a substantial fraction of equity after the IPO. We explain in Section 4.2. that our model can be extended to make it optimal for the entrepreneur to retain some equity after the IPO, and that all of our conclusions obtain also in this extended (and more complicated) model.

stocks and bonds. We can then restate condition (16) as $M_\tau \geq P_\tau$. That is, an IPO takes place if and only if the firm's market value exceeds the firm's private value.

Proposition 1: An IPO at time τ is more likely when

- (a) the benefits of private control, α , are lower
- (b) the uncertainty about average profitability, $\hat{\sigma}_\tau$, is higher
- (c) the idiosyncratic component of the volatility of profitability, $\sigma_{\rho,2}$, is higher
- (d) the current and/or expected profitability, ρ_τ and $\hat{\rho}_\tau$, are higher

Part (a) follows since benefits of private control are consumed only when the firm is private. Lower benefits of private control make an IPO more likely by reducing the opportunity cost of going public. The intuition behind parts (b) and (c) is also simple. If the firm is privately owned, higher uncertainty $\hat{\sigma}_\tau$ or idiosyncratic volatility $\sigma_{\rho,2}$ make the entrepreneur's future consumption more volatile. The risk-averse entrepreneur dislikes this volatility because he is not diversified, and the only way he can diversify is by selling the firm in an IPO. Since investors are well diversified, they are in a better position to bear the risk associated with the private firm's cash flow process. For plausible parameter values, part (c) holds not only for idiosyncratic volatility but also for total volatility $\sigma_\rho \sigma'_\rho = \sigma_{\rho,1}^2 + \sigma_{\rho,2}^2$.

Part (d) follows from the fact that the market value of the firm increases with ρ_τ and $\hat{\rho}_\tau$ more rapidly than the private value does. The effect of expected future profitability, $\hat{\rho}_\tau$, is stronger and easier to explain. Recall from equation (13) that the entrepreneur's consumption is derived from assets in place (B_t) rather than from future growth opportunities. The firm's private value is therefore less sensitive to $\hat{\rho}_\tau$ than the firm's more forward-looking market value is. Increases in $\hat{\rho}_\tau$ push up the private value (because this value is derived from B_t , which grows at the rate of ρ_t , and this growth rate is now expected to be higher) but they push up the market value even more (because the higher expected future profitability is immediately capitalized in the market value). Therefore, higher $\hat{\rho}_\tau$ makes an IPO more likely: The entrepreneur is more willing to forego private benefits in exchange for financial wealth, because doing so moves him to a more valuable consumption path.

The new consumption path is more valuable in part because it is smoother over the entrepreneur's lifetime. When $\hat{\rho}_\tau$ increases, the entrepreneur expects higher consumption in the future. He wants to smooth his consumption by consuming more today but he cannot; his consumption is given in equation (13). (This seems realistic; entrepreneurs cannot easily borrow against expected future profitability.) If $\hat{\rho}_\tau$ is sufficiently high, the entrepreneur's consumption path under private ownership becomes so unattractively steep that he prefers

to sell the firm. After cashing out in an IPO, the entrepreneur increases his consumption immediately and smoothes his consumption path by trading stocks and bonds.

Proposition 2: An IPO takes place at time τ if and only if

$$\hat{\rho}_\tau \geq \underline{\rho}(x_\tau; \hat{\sigma}_\tau, \sigma_\rho), \quad (17)$$

where $\underline{\rho}(x_\tau; \hat{\sigma}_\tau, \sigma_\rho)$ is defined in equation (36) in the Appendix and $x_\tau = \rho_\tau - \hat{\rho}_\tau$.

In words, an IPO takes place if expected profitability is sufficiently high. This rule is similar to the cutoff rule assumed in the toy model in Section 2., except that the cutoff $\underline{\rho}(x_\tau; \hat{\sigma}_\tau, \sigma_\rho)$ here is endogenous: it depends on the model parameters including uncertainty and volatility, and it is also decreasing in x_τ . Proposition 2 is obtained by modifying the IPO condition (16). The intuition behind Proposition 2 is the same as that behind Proposition 1(d). When $\hat{\rho}_\tau$ rises, the market value rises faster than the private value because the former value is more sensitive to $\hat{\rho}_\tau$. When $\hat{\rho}_\tau$ rises sufficiently, an IPO is optimal.

There may be other models that deliver an IPO cutoff rule like ours. For example, consider an asymmetric-information model a la Leland and Pyle (1977) in which an entrepreneur seeking IPO financing must signal high effort to outside investors. It seems plausible for high average profitability to serve as a signal of high effort, which could make an IPO optimal if and only if average profitability exceeds a cutoff. Therefore, we believe that our learning mechanism could also be embedded in asymmetric-information models, with the same basic implications. Our primary interest is in the learning-based implications of the cutoff rule, however this rule is rationalized, for firm profitability around the IPO.

4. Profitability Dynamics Around an IPO

In this section, we analyze the dynamics of profitability around an IPO. Without conditioning on an IPO, profitability ρ_t follows the simple mean-reverting process in equation (10), and expected profitability $\hat{\rho}_t$ is a martingale (equation (27) in the Appendix). Conditioning on an IPO changes the dynamics of ρ_t and $\hat{\rho}_t$ in an interesting way, as we show below.

First, we choose the baseline parameter values to calibrate the model. We take the parameters for the profitability process, $\sigma_{\rho,1} = 0.0584$, $\sigma_{\rho,2} = 0.0596$, and $\phi = 0.3968$ from Pástor and Veronesi (2003) who estimate them from the firm-level return on equity data in 1962–2000. We also choose the same horizon $T = 15$ years, the same pricing kernel volatility $\sigma_\pi = 0.6$, and the same risk-free rate $r = 0.03$ per year as Pástor and Veronesi. These authors report the grand median of profitability of 0.11 per year for public firms. For a typical private firm, the average profitability $\bar{\rho}$ should be lower than 0.11 because only private firms whose

average profitability is perceived to be sufficiently high go public in the model. Therefore, we choose a lower prior mean of $\bar{\rho}$, $\hat{\rho}_0 = 0.07$. We set the prior uncertainty equal to $\hat{\sigma}_0 = 0.05$, so the two-standard-deviation prior bounds for $\bar{\rho}$ are -0.03 and 0.17 per year. We pick $\tau = 5$ years, which is close to the median age of IPO firms in the 1990s (Loughran and Ritter, 2004). We choose risk aversion $\gamma = 2$, the subjective discount rate $\beta = 0.03$, and $\eta = 1$. We consider two values of initial profitability, $\rho_0 = \hat{\rho}_0 = 0.07$ and $\rho_0 = 0$. The latter choice is motivated by the fact that private firms typically do not produce profits when they are started. Measuring the benefits of private control is difficult. We choose $\alpha = 0.10$, a round number.¹¹ Later on, we analyze the sensitivity of our results to α and we also average across many plausible values of α when analyzing the expected post-IPO drop in profitability.

We simulate 10,000 paths of shocks from the model and average the profitability paths across all simulations in which it is optimal for an IPO to take place. First, we draw $\bar{\rho}$ from its prior distribution in equation (14). Starting from ρ_0 , we simulate the realizations of ρ_t between times 0 and T by discretizing the process (10). Analogously, we simulate π_t from the process (11). Given the series of ρ_t and π_t , we compute the dynamics of the posterior beliefs about $\bar{\rho}$. We then check whether the IPO condition (16) is satisfied at time τ . If it is, we keep the simulated path; otherwise we discard it. In the specifications considered below, the IPO condition is satisfied in a little over 60% of all simulated paths.

Figure 2 plots the average paths of realized profitability (ρ_t ; solid line) and expected profitability ($\hat{\rho}_t$; dashed line), where the averages are computed across the simulated paths in which an IPO takes place at time $\tau = 5$. Given the large number of simulations, these average paths represent the model-implied expected patterns in ρ_t and $\hat{\rho}_t$ conditional on an IPO. In Panel A, the initial profitability $\rho_0 = \hat{\rho}_0$; in Panel B, $\rho_0 = 0$. In both panels, realized profitability ρ_t rises sharply before the IPO and declines after the IPO, on average. Expected profitability $\hat{\rho}_t$ also rises before the IPO but it remains flat after the IPO.

To understand the pattern in expected profitability, $\hat{\rho}_t$, recall that for an IPO to take place at time τ , $\hat{\rho}_\tau$ must exceed $\underline{\rho}$. Ex ante, $\hat{\rho}_t$ is a martingale, but the ex-post conditioning on $\hat{\rho}_\tau \geq \underline{\rho}$ implies that $\hat{\rho}_t$ is expected to increase before the IPO. Indeed, in Figure 2, $\hat{\rho}_t$ rises from 0.07 to almost 0.09 between times 0 and τ . After the IPO, there is no conditioning on an ex post event, so $\hat{\rho}_t$ is constant in expectation due to its martingale property.

The pattern in realized profitability, ρ_t , is also intuitive. Recall that expected profitability $\hat{\rho}_t$ must increase before the IPO. An expectation is revised upward only if the realization is higher than expected. To cause upward revisions in $\hat{\rho}_t$, ρ_t must repeatedly be higher than

¹¹Benninga et al (2005) use a range of private benefits centered on 10% of cash flow in their simulations.

expected, so ρ_t rises before the IPO, on average. In order to “pull $\hat{\rho}_t$ up” via Bayesian updating, ρ_t generally rises above $\hat{\rho}_t$, so realized profitability at the time of the IPO is expected to exceed its perceived long-run mean: $\rho_\tau > \hat{\rho}_\tau$. Since $\hat{\rho}_t$ has no post-IPO drift, $\rho_\tau > \hat{\rho}_\tau$ implies that ρ_t is expected to fall after time τ . If the current draw of a random variable (ρ_τ) is above its mean ($\hat{\rho}_\tau$), we expect the future draws to be smaller.

The rise and fall in profitability around the IPO follow from Bayesian learning and the endogeneity of the IPO decision. This pattern does not hinge on the assumption that ρ_t follows the mean-reverting process in equation (10); after all, the same result obtains in the toy model in Section 2., in which profitability is i.i.d. We assume mean reversion in the full model because this assumption is realistic and we want our calibration to be realistic, but our key predictions obtain also in the absence of mean reversion in profitability.

Mean reversion per se can deliver a similar pattern in ρ_t in the absence of learning. The case of no learning is a special case of our model in which $\bar{\rho}$ is a known constant, so that $\hat{\rho}_t = \bar{\rho}$ and $\hat{\sigma}_t = 0$ for all t . In that case, an IPO takes place at time τ if and only if ρ_τ exceeds a cutoff. Related mean-reversion arguments have been proposed by Degeorge and Zeckhauser (1993) for reverse LBOs and by Li, Livdan, and Zhang (2006) for SEOs.

We emphasize that the mean reversion mechanism is distinct from the learning mechanism. Either mechanism can generate a post-IPO decline in profitability. In our toy model, only the learning mechanism is present; in our full model, both mechanisms coexist. We focus on the learning mechanism, which is new. Conveniently, learning generates a richer set of predictions than mean reversion; e.g., mean reversion does not predict a larger post-IPO drop in profitability for firms with lower uncertainty about average profitability.

Figure 3 shows the sensitivity of the profitability pattern in Figure 2 to changes in the baseline parameter values. We change one parameter at a time. The solid line plots the baseline case. The dotted line plots ρ_t for a lower value of prior uncertainty, $\hat{\sigma}_0 = 0.04$. The post-IPO fall in ρ_t is larger than in the baseline case. When uncertainty is lower, prior beliefs about $\bar{\rho}$ are stronger, so ρ_t must rise higher relative to $\hat{\rho}_t$ in order to pull $\hat{\rho}_t$ above any given IPO cutoff. The dashed line plots ρ_t for more volatile profitability, which we obtain by increasing both $\sigma_{\rho,1}$ and $\sigma_{\rho,2}$ to 0.065. The rise and fall in ρ_t are steeper than in the baseline case. Higher volatility makes ρ_t a less precise signal about $\bar{\rho}$, so ρ_t must rise higher relative to $\hat{\rho}_t$ in order to pull $\hat{\rho}_t$ above a given IPO cutoff. When either uncertainty or volatility are lower, the IPO cutoff is generally higher because private ownership is more valuable to the entrepreneur (Proposition 1). But this dependence of the cutoff on volatility and uncertainty does not change the basic predictions of the learning mechanism, as shown in Figure 3.

4.1. Expected Drop in Profitability After an IPO

We have a closed-form solution for the expected long-run post-IPO drop in profitability.

Proposition 3: At time $t < \tau$, the expected post-IPO drop in profitability is given by

$$E_t [\rho_\tau - \hat{\rho}_\tau | \text{IPO at } \tau] = \frac{e^{-\phi(\tau-t)} x_t - \int x_\tau \mathcal{N}(k(x_\tau, \tau; t, x_t, \hat{\rho}_t, \hat{\sigma}_t^2)) \Phi(x_\tau; \mu_x, \sigma_x^2) dx_\tau}{1 - \int \mathcal{N}(k(x_\tau, \tau; t, x_t, \hat{\rho}_t, \hat{\sigma}_t^2)) \Phi(x_\tau; \mu_x, \sigma_x^2) dx_\tau} \quad (18)$$

where $\mathcal{N}(\cdot)$ is the cumulative density function of the standard normal distribution and $\Phi(\cdot; \mu_x, \sigma_x^2)$ is the probability density function of the normal distribution with mean μ_x and variance σ_x^2 . The formulas for $k(\cdot)$, μ_x , and σ_x^2 are given in the Appendix.

We compute the expected drop as of time $t = 0$ for a wide range of parameter values. We vary uncertainty $\hat{\sigma}_0$ from 0 to 10% per year, and both components of volatility, $\sigma_{\rho,1} = \sigma_{\rho,2}$, from 1% to 10% per year. We take $\rho_0 = 0$ and average the results across a range of values for α and $\hat{\rho}_0$: We assume that α is uniformly distributed in [5%, 15%] and $\hat{\rho}_0$ is uniformly distributed in [-20%, 40%]. The remaining parameters are at their baseline values.

It is important here to incorporate the endogeneity of the private firm's existence. We have assumed thus far that the private firm exists but it is not guaranteed that the entrepreneur will start the firm at time 0. Starting the firm is optimal if and only if the condition (40) in the Appendix is satisfied. The parameter sets for which (40) is not satisfied are irrelevant for analyzing IPO profitability; firms with such parameters never go public because they are never created in the first place. This endogeneity can affect the expected post-IPO drop. For example, this drop is typically lower if private benefits are lower (because the IPO cutoff is lower). However, if private benefits are too low, it is not optimal for the entrepreneur to start a private firm at time 0. Therefore, private firms characterized by very low benefits of private control do not exist, and the fact that the post-IPO drop would be low for such firms is nothing more than an intellectual curiosity. To account for the endogeneity of the private firm's existence, we average the expected post-IPO drops only across those sets of parameters for which it is optimal to start the firm at time 0.

Table 1 shows the results. The entries in Panel A range from -0.99% to 23.49% per year, confirming that the expected post-IPO drop in profitability is generally positive. The expected drop is negative only when profitability exhibits unrealistically little volatility. When volatility is low, signals are precise, so learning is fast and $\hat{\rho}_t$ rises rapidly toward the IPO cutoff. Realized profitability ρ_t , which is initiated at $\rho_0 = 0$, may not "catch up" with $\hat{\rho}_t$, in which case we have $\rho_\tau < \hat{\rho}_\tau$ at time τ , after which we expect an increase in profitability. For plausible values of volatility, though, the model-implied magnitudes of the expected drop are not only positive but also close to their empirical counterparts.

Panel A of Table 1 also shows that the expected drop tends to be high when volatility is high and when uncertainty is low. The effects are not fully monotonic due to the endogeneity of the private firm's existence. For example, when uncertainty is high, the firm is less likely to be created at time 0. The firms that are created tend to compensate for the high uncertainty with high values of α , for which the drop is larger. Despite this complicated firm selection effect, the basic patterns in Panel A confirm the implications of the toy model.

In addition to some parameter sets being inadmissible due to failing the condition (40), other sets seem implausible because they imply unrealistic properties for the firm's stock returns. Panels B and C of Table 1 report the average volatility and the average expected excess return on the firm's stock. Both averages are computed in the same way as in Panel A. Stock returns look reasonable when $\sigma_{\rho,1} = \sigma_{\rho,2} > 3\%$, with return volatility ranging from 14% to 45% per year and the expected excess return ranging from 5.9% to 14.8% per year. However, lower values of the profit volatility seem implausible. For example, for $\sigma_{\rho,1} = \sigma_{\rho,2} = 1\%$, return volatility ranges from 3.5% to only 6.6% and the expected excess return is only 1.5%. Since the expected drop in Panel A is non-positive only for the lowest values of the profit volatility, this additional return-based evidence strengthens the conclusion that the expected drop is positive in this model.

4.2. Model Extensions

The model makes several strong assumptions which allow us to obtain analytical solutions. In this section, we relax some of these assumptions. We allow the entrepreneur to

- choose the IPO time optimally (Section 4.2.1.)
- consume and invest optimally before the IPO (Section 4.2.2.)
- invest some of the IPO proceeds back into the firm (Section 4.2.2.)
- sell only a part of the firm in the IPO (Section 4.2.2.)

In all cases, we solve the model numerically, and show that the model's predictions for the dynamics of profitability are robust to these realistic modifications.

4.2.1. Optimal IPO Time

In this section, we relax the assumption that the IPO time τ is exogenously given. The entrepreneur solves for the optimal time τ^* , $0 \leq \tau^* \leq T$, to exercise his option to go public. No IPO is a possibility. The entrepreneur chooses τ^* to maximize lifetime expected utility:

$$\mathcal{V}(B_t, \rho_t, \hat{\rho}_t, t) = \max_{\tau^*} E_t \left[\int_t^{\tau^*} e^{-\beta(s-t)} \frac{c_s^{1-\gamma}}{1-\gamma} ds + e^{-\beta(\tau^*-t)} V(M_{\tau^*}, \tau^*) \right]. \quad (19)$$

We compute τ^* by solving the Bellman equation implied by (19). The details are in the Appendix. We simulate 100,000 paths of shocks and solve for the optimal τ^* for each path.

Figure 4 is the analog of Figures 2 and 3 when τ^* is chosen optimally. To produce a meaningful comparison with figures in which τ is fixed at 5 years, Figure 4 reports average profitability across all simulations in which τ^* is in a one-year window centered at $\tau = 5$ years (i.e., $4.5 \leq \tau^* \leq 5.5$). This event occurs in 5.4% of the simulated paths in Panels A and B (where $\rho_0 = 7\%$), and in 6% of the paths in Panels C and D (where $\rho_0 = 0$).

Figure 4 shows that the model's key predictions are unaffected by making τ endogenous. Profitability rises before the IPO and falls after the IPO, on average, and this pattern is stronger when volatility is high and when uncertainty is low. The rise and fall in profitability are even larger than in Figures 2 and 3. The only notable difference from the $\tau = 5$ case is that prior to its pre-IPO runup, profitability in Panels A and B declines slightly in the first two years of the private firm's existence. The reason is that some of the firms whose early profitability exceeds expectations go public well before year 5. Firms that go public within one year of $\tau = 5$ tend to have stellar profits in years 3 through 5 (otherwise they would not go public around year 5), but on average, they have slightly disappointing profits in years 1 and 2 (otherwise some of them would go public well before year 5). Therefore, our model implies that profitability should rise shortly before the IPO but not necessarily from the beginning of the private firm's existence, as one might infer from Figures 2 and 3.

For an IPO to occur at the optimal time τ^* , $\hat{\rho}_t$ must rise shortly before τ^* , so ρ_t must also rise unexpectedly shortly before τ^* . In the model with an exogenous IPO time τ , the same result holds on average but not without exception because $\hat{\rho}_\tau$ may be high due to high realizations of ρ_t long before τ . In such cases, it might be optimal for an IPO to occur long before the exogenously given τ . In that sense, the intuition behind the profitability patterns comes across even more cleanly when τ is endogenous. Alas, the model with endogenous τ is a bit of a black box because no closed-form solutions (or propositions) are available. By focusing on exogenous τ , we can characterize the model's solution in much more detail.

4.2.2. Other Extensions

We now briefly describe the results from additional extensions of the model. We do not tabulate these results, to save space, but we can make them available upon request.

First, we relax the assumption that the entrepreneur's pre-IPO consumption is given by private benefits αB_t (equation (13)). Instead of forcing the entrepreneur to reinvest all earnings in the firm, we allow him to invest and consume optimally (i.e., consume the optimal

fraction of earnings in each period and reinvest the rest in the firm). We still need to let the entrepreneur consume private benefits. Otherwise, he would never find it optimal to start the private firm at time 0, as there would be no benefit to offset his underdiversification. We model the consumption from private benefits as proportional to the consumption from earnings. We find that the optimal consumption is of the form $c_t = \bar{\alpha}(\rho_t, \hat{\rho}_t, t)B_t$, which is the same form as in equation (13), except that $\bar{\alpha}$ increases with ρ_t and $\hat{\rho}_t$. As a result, an increase in $\hat{\rho}_t$ increases not only future consumption (by increasing the growth rate of B_t) but also current consumption. Since the entrepreneur prefers smooth consumption, $\bar{\alpha}$ is not highly sensitive to $\hat{\rho}_t$, and so the firm's private value remains less sensitive than the market value to $\hat{\rho}_t$. We calibrate the model so the private benefits account for one third of the entrepreneur's pre-IPO consumption. We find that the profitability patterns in this extended model are the same as in Figures 2, 3, and 4.

Second, we relax the assumption that all of the IPO proceeds are invested in stocks and bonds. In reality, some of the IPO proceeds are invested in the firm. We extend the model to allow the entrepreneur to raise additional capital in the IPO and invest it in the firm. We calibrate the model so that the size of the firm doubles after the IPO (i.e., the additional amount of capital raised in the IPO equals the pre-IPO size of the firm). Again, we find the same model-implied profitability patterns as in Figures 2, 3, and 4.

Finally, we let the entrepreneur sell only a part of the firm in the IPO. In the model, selling the whole firm is optimal, but in reality, the entrepreneur usually retains some equity. We accommodate this fact in two ways. First, we force the entrepreneur to (suboptimally) retain a positive fraction of equity. The entrepreneur optimally responds by reducing his post-IPO allocation to stocks. For any given retention ratio, we find the same profitability patterns as in Figures 2, 3, and 4. Second, we extend the model to allow the private benefits to flow also after the IPO, in proportion to the stake in the firm (e.g., if the entrepreneur retains half of the equity, he continues receiving half of the private benefits). In this setting, we can solve for the optimal retention ratio. Conditional on this optimal choice, once again, we find the same model-implied profitability patterns as in Figures 2, 3, and 4.

5. Empirical Analysis

In this section we test the three main predictions of our model: Firm profitability drops after the IPO, on average, and this decline is larger for firms with more volatile profitability and firms with lower uncertainty about average profitability.

5.1. Data

Our data sources include CRSP, Compustat, IBES, SDC, and Jay Ritter’s IPO database. Our sample contains 7,183 firms that had IPOs in the U.S. in 1975–2004. We include a firm in the sample if it meets all of the following criteria: (1) it appears in either Jay Ritter’s 1975–1984 IPO database or in SDC’s New Issues database with an offer date between 1/1/1985 and 12/31/2004; (2) it had a firm-commitment IPO; (3) it is not a closed-end fund, trust, unit, ADR, ADS, or REIT; and (4) the IPO’s offer price was at least \$1 per share.

We measure profitability as earnings scaled by the book value of equity, or return on equity (ROE). $ROE_{i,s}$ is computed for firm i in the fiscal quarter that is s quarters after the IPO. The dependent variable in our tests is $ROE_{i,s} - ROE_{i,0}$, the change in ROE over the first s quarters after firm i ’s IPO. ROE equals earnings divided by book equity. We calculate earnings using quarterly Compustat data, and book equity using both quarterly and annual Compustat data. Further details on the construction of $ROE_{i,s}$ are in the Appendix.

We estimate the volatility of ROE by the standard deviation of quarterly ROE over a five-year period after the IPO. Specifically, $VOL(i; s_0)$, or $VOL(s_0)$ for short, is the standard deviation of $ROE_{i,s}$ in quarters $s = s_0, \dots, s_0 + 19$, assuming that at least 12 observations are available. We use two values of s_0 . The natural choice is $s_0 = 0$ because $VOL(0)$ uses data as close to the IPO as possible. Under this choice, some of the earnings data used to compute $VOL(0)$ are also used to compute the dependent variable, $ROE_{i,s} - ROE_{i,0}$. Firms with large post-IPO increases or decreases in ROE are likely to have large values of $VOL(0)$. Although there is no obvious bias, we address this concern by using $s_0 = s + 1$. There is no overlap between the earnings data used to calculate $VOL(s + 1)$ and $ROE_{i,s} - ROE_{i,0}$.

5.2. Separating Uncertainty from Volatility

Commonly used proxies for uncertainty such as firm age, size, return volatility, or analyst coverage are inadequate here because they proxy not only for uncertainty but also for the volatility of profitability, which has an opposite theoretical effect on the post-IPO drop in profitability. Firms with high uncertainty also tend to have high volatility, which presents an estimation challenge. Luckily, we have found an empirical proxy whose value should be high when uncertainty is high and when volatility is low: the stock price reaction to earnings announcements. In fact, we can link this proxy directly to our model.

Corollary 1: If the model’s assumptions hold and, in addition, $\sigma_{\rho,1} = 0$, then

$$dR_t - E_t[dR_t] = M(\sigma_{\rho,2}, \hat{\sigma}_0^2; \phi, t)(d\rho_t - E_t[d\rho_t]), \quad (20)$$

where

$$M(\sigma_{\rho,2}, \hat{\sigma}_0^2; \phi, t) = Q_1(T-t) + Q_2(T-t)\phi \frac{\hat{\sigma}_t^2}{\sigma_{\rho,2}^2}, \quad (21)$$

$dR_t - E_t[dR_t]$ is the unexpected stock return, and Q_1 and Q_2 are in the Appendix.

The quantity M represents the stock price reaction to earnings surprises, typically called the earnings response coefficient (e.g., Easton and Zmijewski, 1989). M is positive in the model because earnings surprises and the associated abnormal returns have the same sign. M is increasing in uncertainty ($\hat{\sigma}_t$) and decreasing in volatility ($\sigma_{\rho,2}$). The intuition is clear. Realized earnings are a noisy signal about average future profitability. Upon observing a given signal, investors update their beliefs about the firm value more when they are more uncertain and when the signal is less noisy (i.e., when earnings are less volatile).¹²

Our model predicts that firms with higher values of M have smaller post-IPO drops in profitability, because such firms have higher uncertainty, lower volatility, or both (holding ϕ and t constant). Once we control for profit volatility, the regression of $ROE_{i,s} - ROE_{i,0}$ on M_i can be interpreted as a test of the model's prediction regarding uncertainty. The theoretical motivation for M is only approximate because Corollary 1 requires $\sigma_{\rho,1} = 0$. We expect this approximation to be reasonably good, though, because we estimate M in short periods around firm-level earnings announcements, during which firm-specific earnings news is the main driver of unexpected stock returns. While M is not a perfect proxy, it cleanly separates uncertainty from volatility, and it is also directly motivated by the model.

We estimate M_i for each IPO firm i based on post-IPO earnings announcement data.¹³ On the left-hand side of equation (20), we interpret $dR_t - E_t[dR_t]$ as the abnormal return due to an earnings announcement. We measure this quantity by AR_{it} , the cumulative return of stock i in excess of stock i 's industry's return starting one trading day before the firm's t -th post-IPO earnings announcement and ending one trading day after the same announcement. Quarterly earnings announcement dates are from IBES. Daily stock returns are from CRSP, and daily returns of 49 value-weighted industry portfolios are from Ken French's website. On the right-hand side of equation (20), we interpret $d\rho_t - E_t[d\rho_t]$ as unexpected quarterly profitability, which we compute as $(EPS_{it} - E[EPS_{it}]) / BE_{it}$. EPS_{it} denotes the quarterly earnings per share of firm i announced in its t -th post-IPO earnings announcement, from

¹²Half of this intuition is already in Lang (1991), who argues that M should be larger when there is more uncertainty about the firm's earnings process. Consistent with his argument, Lang finds that the estimates of M tend to decrease as a firm gets older (and more earnings data become available). Lang's evidence strengthens the motivation for our use of M as a proxy for uncertainty after controlling for volatility.

¹³To estimate uncertainty and volatility at the time of the IPO, we use post-IPO data, for several reasons. First, we do not have pre-IPO data on stock returns and earnings. Second, in the model, profit volatility is the same before and after the IPO. Finally, firms with higher pre-IPO uncertainty also generally have higher post-IPO uncertainty, because uncertainty in the model declines deterministically over time, both before and after the IPO, due to learning. The IPO itself has no effect on uncertainty.

the IBES unadjusted actuals file. $E[EPS_{it}]$ is the mean of all analyst forecasts of EPS_{it} using IBES's last pre-announcement set of forecasts for the given fiscal quarter. BE_{it} is book equity per share of firm i , using the most recent pre-announcement measurement.

To estimate M_i , we compute two measures, $ERC_1(i)$ and $ERC_2(i)$. First, we compute

$$RC_{it} = \frac{AR_{it}}{(EPS_{it} - E[EPS_{it}]) / BE_{it}}, \quad (22)$$

which is a proxy for M_i (see equation 20). Since RC_{it} is noisy (especially if the denominator is close to zero), we winsorize the highest 5% and lowest 5% of RC_{it} observations. We also average the quarterly RC_{it} 's over the first three years after the IPO to increase precision:

$$ERC_1(i) = \frac{1}{13} \sum_{t=0}^{12} RC_{it}. \quad (23)$$

We compute $ERC_1(i)$ only if there are at least six valid observations of RC_{it} . To define $ERC_2(i)$, consider the following regression over the five-year period after the IPO:

$$(EPS_{it} - E[EPS_{it}]) / BE_{it} = \gamma_{i0} + \gamma_{i1} AR_{it} + \varepsilon_{it}, \quad t = 0, 1, \dots, 20. \quad (24)$$

According to equation (20), $\gamma_{i1} = 1/M_i$ but we do not measure M_i as $1/\hat{\gamma}_{i1}$ because $\hat{\gamma}_{i1}$ can be close to zero, producing outliers in $1/\hat{\gamma}_{i1}$. Instead, we define

$$ERC_2(i) = -\hat{\gamma}_{i1}, \quad (25)$$

with a minus sign so that large earnings responses are associated with large values of ERC_2 . Unlike ERC_1 , ERC_2 is not a direct estimate of M , but it preserves the same cross-sectional ranking. We make earnings surprises the dependent variable in equation (24) to mitigate the attenuation bias, since we believe there is more measurement error in earnings surprises than in abnormal returns.¹⁴ Since equation (20) indicates $\gamma_{i0} = 0$, we estimate the regressions in (24) without the intercept. We require at least 10 observations to estimate these regressions. Before running the regressions, we winsorize the highest and lowest 5% values of both AR_{it} and $(EPS_{it} - E[EPS_{it}]) / BE_{it}$ across all firms and quarters $t = 0, 1, \dots, 32$.

5.3. Summary Statistics

Table 2 reports some summary statistics. The three-year change in ROE, $ROE_{i,12} - ROE_{i,0}$, can be computed for 4,254 firms. Its mean and median are both negative, as the model predicts. $ROE_{i,12} - ROE_{i,0}$ is negatively correlated with volatility and positively correlated with the ERCs. These correlations foreshadow our main empirical results.

¹⁴Nonetheless, the conclusions regarding the estimates of M_i are the same when we run the regression (24) in reverse; i.e., when we switch the dependent and independent variables in the regression.

Profitability in the quarter of the IPO, $ROE_{i,0}$, can be computed for 6,149 of the 7,183 firms in our sample.¹⁵ The median $ROE_{i,0}$ is 1.69% per quarter. ERC_1 and ERC_2 can be computed for about 40% of firms. (IBES coverage begins only in 1982.) The mean of ERC_1 shows that a 1% earnings surprise (scaled by book equity) is associated with a 3.13% abnormal stock return. Theoretically, earnings surprises and stock returns should have the same sign, so ERC_1 should be positive and ERC_2 negative. Alas, $ERC_1 < 0$ for 33% of firms, and $ERC_2 > 0$ for 22% of firms. These unexpected signs seem due to measurement error in expected earnings and non-earnings related news. The cross-sectional means of ERC_1 and ERC_2 do have the predicted signs and high statistical significance. Since ERC_1 and ERC_2 proxy for uncertainty divided by volatility, we expect them to be negatively correlated with volatility, and they indeed are. Surprisingly, ERC_1 and ERC_2 are almost uncorrelated with each other, but this result is artificially created by the observations of ERC_1 and ERC_2 that do not have the predicted signs.¹⁶ When these observations are excluded, the correlation between ERC_1 and ERC_2 increases substantially. We define ERC_1^+ and ERC_2^- in the same way as ERC_1 and ERC_2 , except that we declare the observations with $ERC_1 < 0$ and $ERC_2 > 0$ as missing. The correlation between ERC_1^+ and ERC_2^- is 0.30.

Figure 5 plots the change in ROE, $ROE_{i,s} - ROE_{i,0}$, in event time following the IPO. The top panel shows that average ROE drops steadily after the IPO, leveling off after about two years. The median change in ROE, shown in the middle panel, is also negative. The 75th percentile line shows that for more than a quarter of firms, ROE increases after the IPO. This is consistent with the model, which makes predictions only about the *average* post-IPO change in ROE (a drop in ROE is expected ex ante but need not happen ex post). The bottom panel shows the mean change in ROE in the sub-samples of firms that had IPOs in 1975–1984, 1985–1994, and 1995–2004. The patterns are remarkably similar across the three sub-samples, and they are also similar to the model-implied pattern in Figure 2.

Figure 6 compares the post-IPO average changes in ROE between firms with high and low values of volatility and the ERCs. We split all firms into two equally large sub-samples based on whether the firms' $VOL(0)$ exceeds the cross-sectional median of $VOL(0)$, and we do the same for ERC_1 .¹⁷ In Panels A and B, we plot the sub-samples' mean changes in ROE. In Panels C and D, we plot the differences between sub-samples, with 95% confidence

¹⁵To mitigate the Compustat backfilling bias described by Chan, Jegadeesh, and Lakonishok (1995), Kothari, Shanken, and Sloan (1995), and others, we do not use pre-IPO accounting data.

¹⁶Under the assumptions that deliver equation (20), ERC_1 and ERC_2 are approximate estimates of M and $-1/M$, respectively, so $ERC_2 \approx -1/ERC_1$. The function $f(x) = -1/x$ is monotonically increasing for $x > 0$ (which is the predicted sign of ERC_1), making x and $f(x)$ perfectly positively correlated, but the presence of negative values of x (i.e., values of ERC_1 with unpredicted signs) destroys this relation since we observe both branches of the hyperbola instead of just the branch with $x > 0$ and $f(x) < 0$.

¹⁷The results based on $VOL(13)$ and ERC_2 lead to the same conclusions.

intervals. Figure 6 shows that mean ROE drops for both high- and low- $VOL(0)$ firms, the drop is significantly larger for firms with high $VOL(0)$, and the difference grows with the horizon. Similarly, mean ROE drops for both high- and low- ERC_1 firms, the drop is larger for low- ERC_1 firms, and the difference generally grows with the horizon. Both results support the model. Since ERC_1 depends on both uncertainty and volatility, it is unclear which of the two variables drives the difference between the high- and low- ERC_1 firms. To disentangle these effects, we include both volatility and the ERCs in a multiple regression.

5.4. Regression Analysis

We estimate the following regression across all IPO firms with available data:

$$ROE_{i,s} - ROE_{i,0} = X_i\beta + \varepsilon_i, \quad (26)$$

where X_i contains a constant and various combinations of our measures of ROE volatility and earnings response (we also include controls in Section 5.5.1.). We consider two horizons, $s = 4$ and $s = 12$ quarters. In each specification, we use as many observations as possible, so the sample is not necessarily the same across specifications. We estimate β by OLS and calculate its standard error by clustering the regression residuals in calendar time.¹⁸

Table 3 shows the results. First, we estimate the unconditional mean change in ROE over the first 4 and 12 post-IPO quarters, respectively. The average value of $ROE_{i,4} - ROE_{i,0}$ is -2.67% per quarter ($t = -10.3$) and the average value of $ROE_{i,12} - ROE_{i,0}$ is -4.19% per quarter ($t = -13.7$). The average post-IPO drop in profitability is consistent with the model as well as with the earlier empirical studies. Moreover, the magnitude of the drop is comparable to the model-implied values reported in Table 1.

Second, we test the model's prediction that ROE drops more for firms with more volatile ROE. Indeed, the slope coefficients on both $VOL(0)$ and $VOL(s+1)$ are negative and highly statistically significant, with t -statistics exceeding 8.6 at both horizons. The relation is also economically significant: a one-standard-deviation increase in $VOL(0)$ is associated with a 1.92 percentage point drop in $ROE_{i,4} - ROE_{i,0}$ and a 5.46% drop in $ROE_{i,12} - ROE_{i,0}$. The corresponding numbers for $VOL(s+1)$ are 1.51% and 2.09%, respectively.

Third, we test the prediction that ROE drops more for firms with smaller earnings response measures. Indeed, the slope coefficients on ERC_1 and ERC_2 are positive in all four

¹⁸We allow non-zero correlations between the residuals of firms whose IPOs were $s/2$ or fewer quarters apart in calendar time. Specifically, we assume that $E[\varepsilon_i\varepsilon_j]$ is equal to σ^2 for $i = j$ and to σ_t^2 for $i \neq j$, where t is the number of quarters between i and j 's IPOs. For $t \leq s/2$, we estimate σ_t^2 from the relevant subset of the estimated OLS residuals; for $t > s/2$, we set $\sigma_t^2 = 0$.

specifications (two horizons, two ERCs), and three of the four coefficients are statistically significant. A one-standard-deviation decrease in ERC_1 is associated with a 0.68 percentage point drop in $ROE_{i,4} - ROE_{i,0}$ and a 1.02% drop in $ROE_{i,12} - ROE_{i,0}$. The corresponding numbers for ERC_2 are 0.35% and 0.60%, respectively.

Fourth, since firms with smaller ERCs should have either lower uncertainty or higher volatility or both, we attempt to isolate the impact of uncertainty by including controls for volatility. In these multiple regressions, the slope coefficients on volatility remain negative and highly significant. The slope coefficients on ERC_1 and ERC_2 are positive in all eight specifications (two horizons, two ERCs, two volatility measures), but only three of these coefficients are statistically significant. These results are consistent with the model's uncertainty prediction, but the evidence is weaker than that for the volatility prediction.

The ERCs may contain substantial estimation error due to mismeasurement of earnings expectations and to non-earnings-related news. This error is likely to affect especially the ERC estimates that do not have the predicted signs (i.e., $ERC_1 < 0$ and $ERC_2 > 0$); in fact, this error is the most likely reason why these signs are opposite to what basic economics would predict. Therefore, we repeat the tests from Table 3 using ERC_1^+ and ERC_2^- .

Table 4 reports the results. First, consider the simple regressions of $ROE_{i,s} - ROE_{i,0}$ on either ERC_1^+ or ERC_2^- . The results show that ROE drops more for firms with smaller ERCs, and the evidence is even stronger than in Table 3: the slope coefficients on ERC_1^+ and ERC_2^- are significantly positive in all four univariate specifications, with t -statistics ranging from 2.90 to 6.67. Second, consider the same regressions but control for the volatility of ROE. The slope coefficients on ERC_1^+ and ERC_2^- are positive in all specifications, and five of the eight coefficients are statistically significant. Again, these results are stronger than in Table 3.¹⁹ This increase in significance suggests that the decrease in precision resulting from a smaller number of observations is more than offset by the increase in precision resulting from using the ERCs that contain less measurement error. These results support the model's prediction that the post-IPO drop in ROE is larger for firms with less uncertainty.

5.5. Robustness Analysis

This section describes additional robustness tests. One set of tests involves modifying the definitions of the variables used in the regressions. First, redefining profitability from ROE

¹⁹We obtain similar results when we winsorize the ERCs with unpredicted signs at zero instead of eliminating them. The slope coefficients on the ERCs are significantly positive in all four univariate specifications, and they are also positive in all eight specifications that control for the volatility of ROE, with four of the eight coefficients being statistically significant.

to ROA (return on assets) makes little difference. ROA drops significantly after the IPO, and this drop is significantly larger for firms with more volatile profits in all specifications. When volatility and ERC are both included in the regression, ERC always has a positive slope. This slope is statistically significant in three specifications, just like in Table 3. When ERC_1 and ERC_2 are replaced by ERC_1^+ and ERC_2^- , the slope on ERC is positive and significant in four out of eight specifications in which volatility is also included, and it is positive and significant in three out of four specifications in which volatility is excluded. In short, the results based on ROA are very similar to those based on ROE.

It makes little difference whether we use the median instead of the mean of analyst forecasts when estimating $E[EPS_{it}]$, or whether we require at least two forecasts to compute the mean. Changing the horizon over which we measure the post-IPO drop in ROE to two years or four years does not change any conclusions. Changing the number of quarters over which ERC_1 and ERC_2 are computed leads to similar results. We obtain very similar results when we free up the intercept in the regression (24), and also when we redefine ERC_2 as the slope in the reverse regression of abnormal returns on earnings surprises. We also include an additional regressor in (24), the cumulative stock return starting one day after IBES records the analyst forecasts and ending two trading days before the earnings announcement, to soak up some of the news that comes out before the earnings announcement but after analysts form their forecasts. The resulting modification of ERC_2 enters our regressions with the same sign but slightly higher standard errors than the original ERC_2 . However, the modified ERC_2 has the predicted sign less often than the original ERC_2 , so including the additional regressor seems to reduce rather than increase precision. Finally, we split the sample by the IPO date into two roughly equally-sized subsamples, 1975–1992 and 1993–2004. In both subsamples, we find results similar to those in the full sample.

5.5.1. Controls

Next, we add various controls on the right-hand side of regression (26). One control whose inclusion seems warranted by the model is the mean reversion coefficient ϕ , which appears in equation (21). We compute sample estimates of ϕ from the profitability series of each firm. We find that ϕ does not enter significantly in regression (26), and that controlling for ϕ leads to exactly the same conclusions as in Section 5.4.

We also include several controls that are not motivated by the model. First, we include firm size, measured by the logarithm of the ratio of the firm’s assets to the total assets of all firms in Compustat at the end of the IPO’s fiscal year. We scale firm size in this way to adjust for inflation between the various IPO dates. Univariate regressions indicate that

profitability drops more for smaller firms, but this relation loses its statistical significance when we control for volatility at the three-year horizon. Second, we include firm investment, measured by averaging the ratios of capital expenditures to total assets over the first four or twelve post-IPO quarters. Univariate regressions show that profitability drops more for firms that invest more after the IPO, but this relation becomes insignificant when we control for either volatility or ERCs. Third, we control for firm leverage, measured by the ratio of total debt to total assets at the end of the IPO quarter. We find that high leverage is associated with a smaller post-IPO drop in ROE over a one-year horizon, but not over a three-year horizon. Most important, including any of the three controls – size, investment, or leverage – has very little effect on the volatility and ERC slopes and no effect on their significance. The same holds when all three controls are included jointly. Our volatility and uncertainty results are therefore robust to controlling for size, investment, and leverage.

Finally, we include industry fixed effects, which control for omitted variables that differ across industries (defined by two-digit SIC codes) but not over time. Including industry fixed effects leads to the same conclusions: All slopes that are statistically significant in Table 4 remain significant. This is not surprising because industry fixed effects explain less than 2% of the variation in the regressand. Table 5 is the counterpart of Table 4 with controls for industry fixed effects, leverage, and size. We drop the specification with just a constant because fixed effects are included. The results are very similar to those in Table 4: All variables that are significant in Table 4 are also significant in Table 5. We conclude that our results are robust to controlling for industry classification, leverage, and size.

5.5.2. Venture Capital Financing

Next, we relate the post-IPO drop in ROE to venture capital (VC) financing.²⁰ We use data from Barry et al (1990), available at Paul Gompers' website, to measure VC involvement in IPOs in 1978–1987. We measure VC involvement (VCI) as the fraction of the firm's equity that is owned by VCs at the time of the IPO. There are 2,192 firms in our sample with IPOs in 1978–1987, 1,061 of which we found in the database of Barry et al (1990). Among these firms, 745 (70%) had no VC involvement. Among the remaining 316 firms, VCI has a mean of 35% and standard deviation of 22%. We find that VCI is positively correlated with $ROE_{i,4} - ROE_{i,0}$ ($t = 2.99$), indicating that ROE drops less for firms with more VC involvement. This result is broadly consistent with our model. There are no VCs in our model but if there were, the diversification motive for an IPO would be weaker because the VCs tend to be better diversified than the entrepreneur. As a result, one might conjecture that the post-IPO drop in ROE would be smaller for firms with higher VCI, and that is

²⁰Lerner (1994) and Black and Gilson (1998), among others, analyze the effect of VCs on the IPO decision.

indeed what we find. The relation between VCI and the change in ROE remains significant at the one-year horizon after controlling for volatility and the ERCs. The relation disappears at the three-year horizon, though, so the support for our VC conjecture is limited.

When we include VCI in regression (26), we find a significantly negative slope on volatility in all specifications, so our volatility result is robust to controlling for VCI. The slopes on the ERCs are mostly positive but insignificant. This result is driven by the small sample rather than by the inclusion of VCI. To compute the ERCs, we need IBES data, and IBES coverage was poor in 1978–1987. As a result, the tests that involve both the ERCs and VCI are based on only 133 to 180 firms, a tiny subset of our overall sample. In this subset, the ERCs do not enter significantly, whether or not we control for VCI. Therefore, we have insufficient data to test whether our uncertainty result is robust to controls for VC involvement.

5.5.3. Earnings Management

Earnings management is a plausible explanation for the post-IPO drop in ROE. If firms inflate their pre-IPO earnings through discretionary accruals, earnings will naturally fall after the IPO. To distinguish the earnings management hypothesis from our learning hypothesis, note that the earnings management hypothesis does not make the same predictions for volatility and uncertainty; in fact, it appears to make the opposite prediction regarding volatility. Firms that are willing to manipulate their earnings around the IPO are likely to manipulate them after the IPO as well. Such firms are likely to smooth their post-IPO earnings, given the apparent market preference for less volatile earnings.²¹ Therefore, earnings management would seem to predict that the post-IPO decline in ROE should be larger for firms with less volatile post-IPO earnings, but we find the opposite result.

We test whether our results are related to earnings management. Following Teoh, Welch, and Wong (1998), we interpret discretionary current accruals (DCA) as the portion of earnings that may have been manipulated, and estimate firms' annual DCA as the fitted residual from a cross-sectional regression of current accruals on the change in sales. In Teoh et al's sample period of 1980–1992, we obtain the same median DCA in the IPO year, 4.0% per year. In our longer sample period, the median DCA in the IPO year is 2.6% per year, and the median DCA in year three after the IPO is zero. The 2.6% drop in median DCA is smaller than the 4.5% drop in median annual ROE over the same horizon, which suggests that earnings management cannot fully explain the post-IPO drop in ROE.²²

²¹For example, Graham, Harvey and Rajgopal (2005) survey 401 financial executives and find that more than three quarters of them would give up economic value in exchange for smooth earnings.

²²There are also other reasons to believe that earnings management cannot fully explain the post-IPO drop in profitability. First, Ball and Shivakumar (2006) argue that IPO firms actually supply more conservative

We include $DCA(0)$, the DCA in the fiscal year of the IPO, on the right-hand side of regression (26). The earnings management hypothesis predicts that firms with larger $DCA(0)$ have larger post-IPO drops in ROE. Although $DCA(0)$ enters with a negative slope in most specifications, this slope is not statistically significant. More important, the inclusion of $DCA(0)$ does not affect the significance of the slope coefficients on volatility and ERC. After controlling for $DCA(0)$, the slopes on volatility remain significantly negative in all specifications, and the slopes on ERC remain positive, with even higher statistical significance. This evidence suggests that our results are robust to controlling for earnings management.

6. Conclusions

This paper develops a dynamic model of the optimal IPO decision, analyzes the model's predictions, and tests these predictions empirically. In the model, two types of agents, well-diversified investors and an under-diversified entrepreneur, learn about the mean profitability of a private firm by observing realized profits. There is no asymmetric information. The entrepreneur faces a tradeoff between benefits of private control and diversification benefits of going public. An IPO takes place when the firm's market value exceeds the private value, which happens when the firm's expected future profitability is sufficiently high. The model predicts that firm profitability declines after the IPO, on average, and that this decline is larger for firms with more volatile profitability and firms with less uncertain average profitability. We test the three predictions empirically and find significant support for them. High volatility and high uncertainty tend to go together, but we separate them by estimating the stock price reaction to earnings announcements, which is strong when uncertainty is high and volatility is low. The model works well not only qualitatively but also quantitatively, in that it can match the magnitude of the post-IPO drop reasonably well.

Our main contribution is a novel learning mechanism that generates a post-IPO drop in profitability without relying on irrationality or asymmetric information. This mechanism also makes new cross-sectional predictions for volatility and uncertainty. These two additional predictions help us distinguish our model from alternatives that do not involve Bayesian learning. While the post-IPO drop in profits can also be in part due to other mechanisms (e.g., asymmetric information, earnings management, mean reversion in profitability, and behavioral stories), none of these alternatives predicts that this drop should be larger for firms with higher volatility and lower uncertainty. Since the volatility and uncertainty predictions

and higher-quality financial reports than other firms. Second, Jain and Kini (1994) show that when earnings are replaced by cash flow (which is more difficult to manipulate) in the numerator of the profitability measure, profitability still drops significantly after the IPO. Finally, Chemmanur et al (2006) find a post-IPO drop in firm-level total factor productivity, which seems immune to accounting manipulation.

seem unique to learning, our empirical evidence in favor of these predictions suggests that learning is at least partly responsible for the observed profitability patterns.

The learning mechanism also seems relevant for the going private decision (e.g., Zingales, 1995, Benninga et al, 2005, Bharath and Dittmar, 2006). Reversing our arguments for going public, a firm is taken private if the benefits of private control exceed the diversification benefits of public ownership, which happens when expected profitability is sufficiently low. Firms should experience declines in profitability before going private and increases in profitability after going private. Consistent with the first prediction, Halpern, Kieschnick, and Rotenberg (1999) find that stock returns before leveraged buyouts are unusually low. Regarding the second prediction, private equity firms often claim to make companies more profitable after taking them private. Our learning mechanism suggests that profitability should rise after going private even without any special skill on behalf of the company's management.

Our model also has some relevance for seasoned equity offerings (SEO). If a shareholder owns a substantial fraction of a firm, she faces a similar tradeoff as our entrepreneur: issuing equity makes her more diversified while reducing her control over the firm. Following the model's logic, the shareholder may find it optimal to issue more equity after profitability rises sufficiently, and profitability should subsequently fall for the same reasons as in the model. Indeed, Loughran and Ritter (1997) find that firm profitability tends to increase before an SEO and decline thereafter, as the model would imply. It would be interesting to test whether this pattern in profitability around SEOs is related to volatility, uncertainty, and to the fraction of equity held by the firm's largest shareholder.

Loughran and Ritter (1997) also argue that "The most salient feature concerning firms' equity issuance behavior is that most firms issue equity after large stock price increases." For example, Asquith and Mullins (1986) and Loughran and Ritter (1995) report that firms engaging in SEOs tend to exhibit high stock returns prior to the SEO. This empirical fact is also consistent with our model. In the model, an issue of equity is induced by recent unexpected increases in profitability, which should coincide with high stock returns. We cannot test this prediction on IPOs since pre-IPO stock returns are obviously unavailable, but the SEO evidence seems comforting. Also note that our model makes no unusual predictions regarding the post-issue stock returns, which are actively debated in the literature.²³ In our model, expected stock returns are always fair, determined by the stock's systematic risk. We focus on operating performance rather than stock performance.

²³For example, Ritter (1991) and Loughran and Ritter (1995) show that stock returns of firms that recently went public are lower on average than returns of seasoned firms, while Brav and Gompers (1997) and Brav, Géczy, and Gompers (2000) argue that most IPOs are small growth stocks and such stocks have had low returns regardless of whether they recently went public.

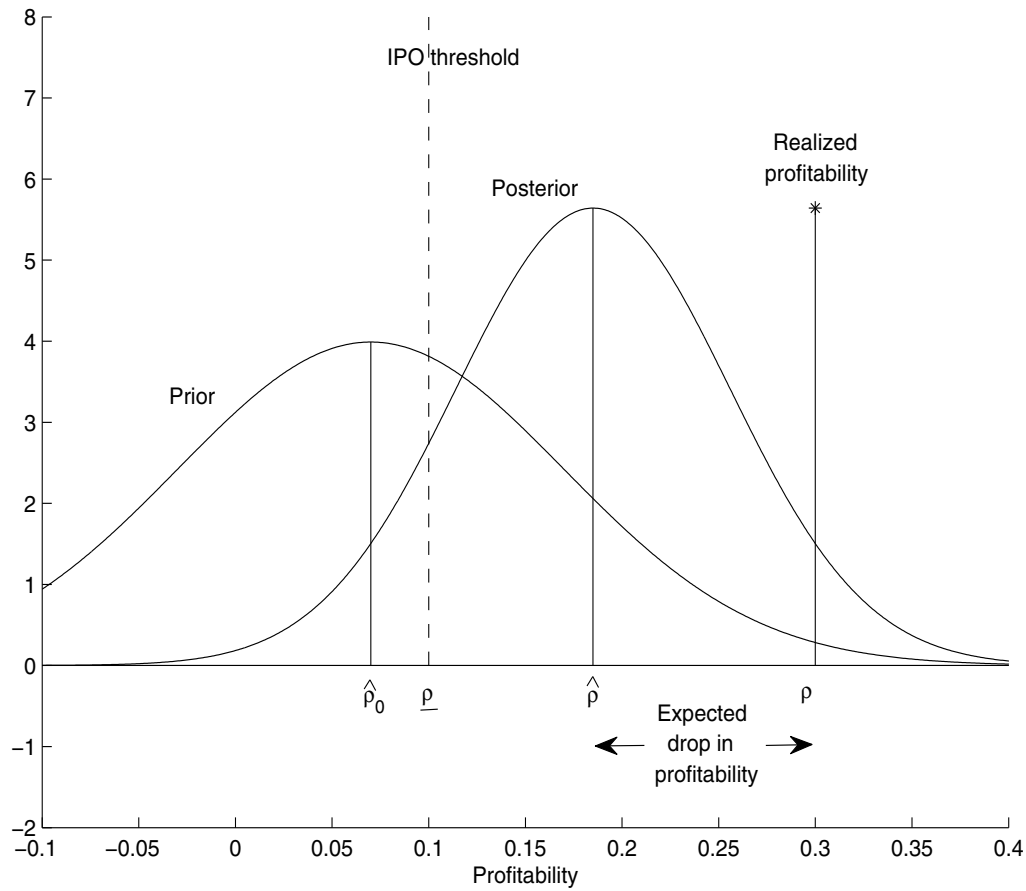


Figure 1. Expected Post-IPO Drop in Profitability in the Toy Model. This figure captures the intuition behind Result 1 from the toy model. An IPO takes place if and only if the expected profitability exceeds the cutoff $\underline{\rho}$. If an IPO takes place at time 1, the expected profitability at time 1 must exceed the cutoff ($\hat{\rho} \geq \underline{\rho}$) but the expected profitability at time 0 must be below the cutoff ($\hat{\rho}_0 < \underline{\rho}$). Since the posterior mean $\hat{\rho}$ is a weighted average of the prior mean $\hat{\rho}_0$ and the signal $\underline{\rho}$, we have $\rho > \hat{\rho}$, and the expected post-IPO drop in profitability, or $E_0(\rho - \hat{\rho} | \text{IPO at time 1})$, is positive.

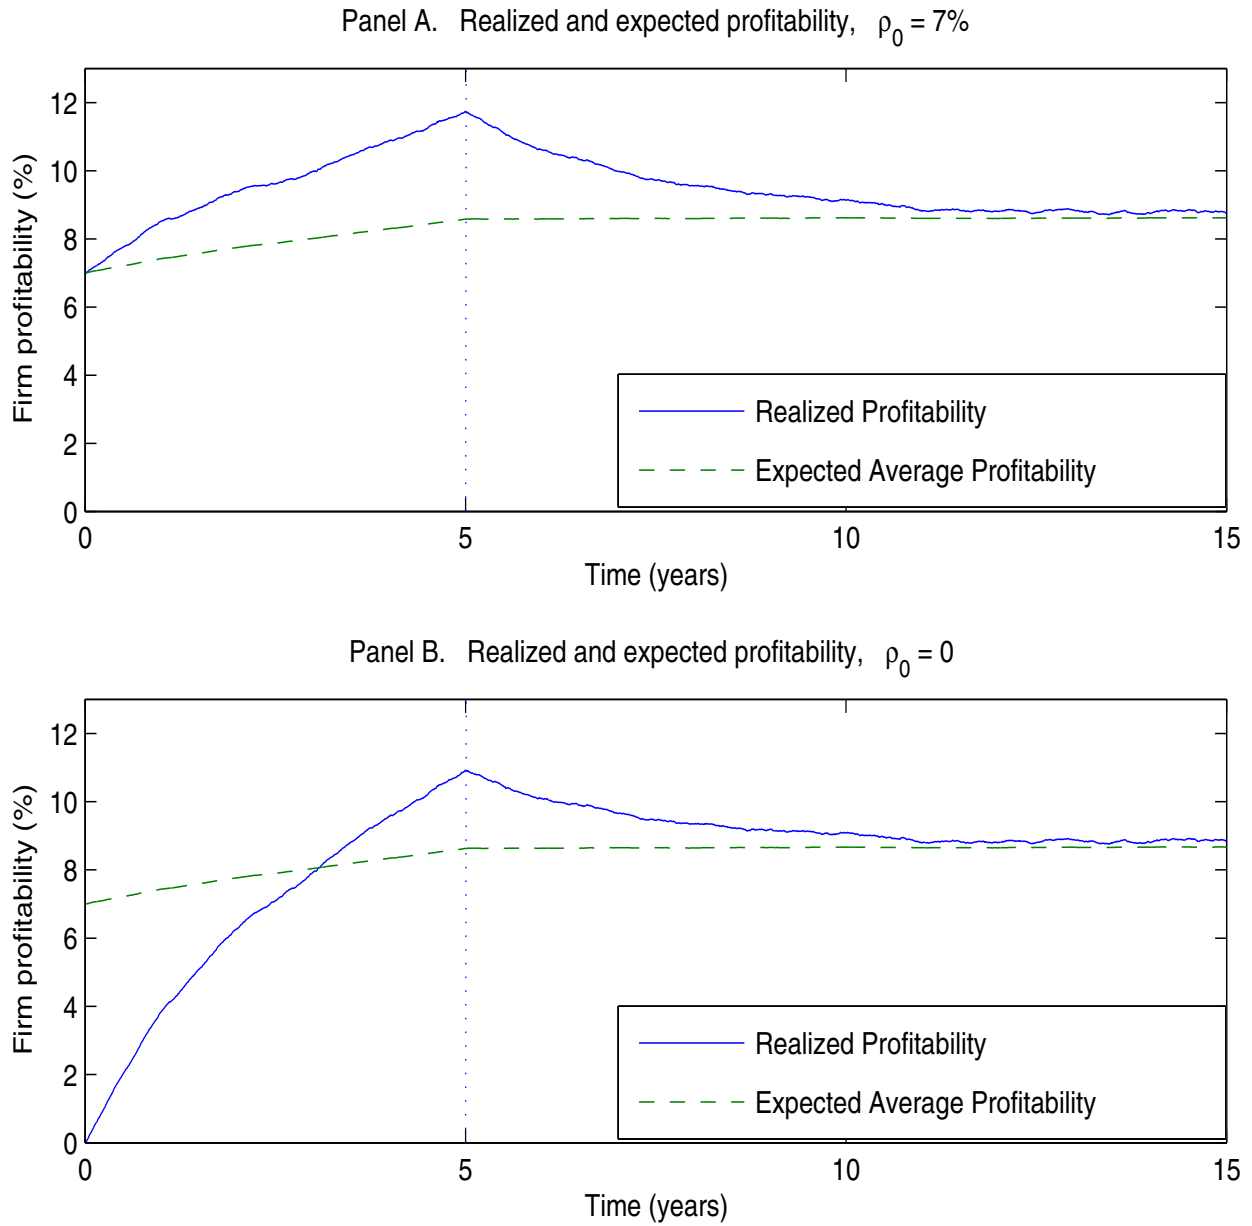


Figure 2. Model-Implied Expected and Realized Profitability Around an IPO.

This figure plots the average paths of realized profitability (ρ_t ; solid line) and expected average profitability ($\hat{\rho}_t$; dashed line), in percent per year, where the paths are averaged across all simulations of our model in which an IPO takes place at time $\tau = 5$. Given the large number of simulations, these average paths represent expected patterns in ρ_t and $\hat{\rho}_t$ conditional on an IPO. Profitability is defined as earnings over book value of equity. In Panel A, the initial profitability $\rho_0 = \hat{\rho}_0 = 7\%$; in Panel B, $\rho_0 = 0$.

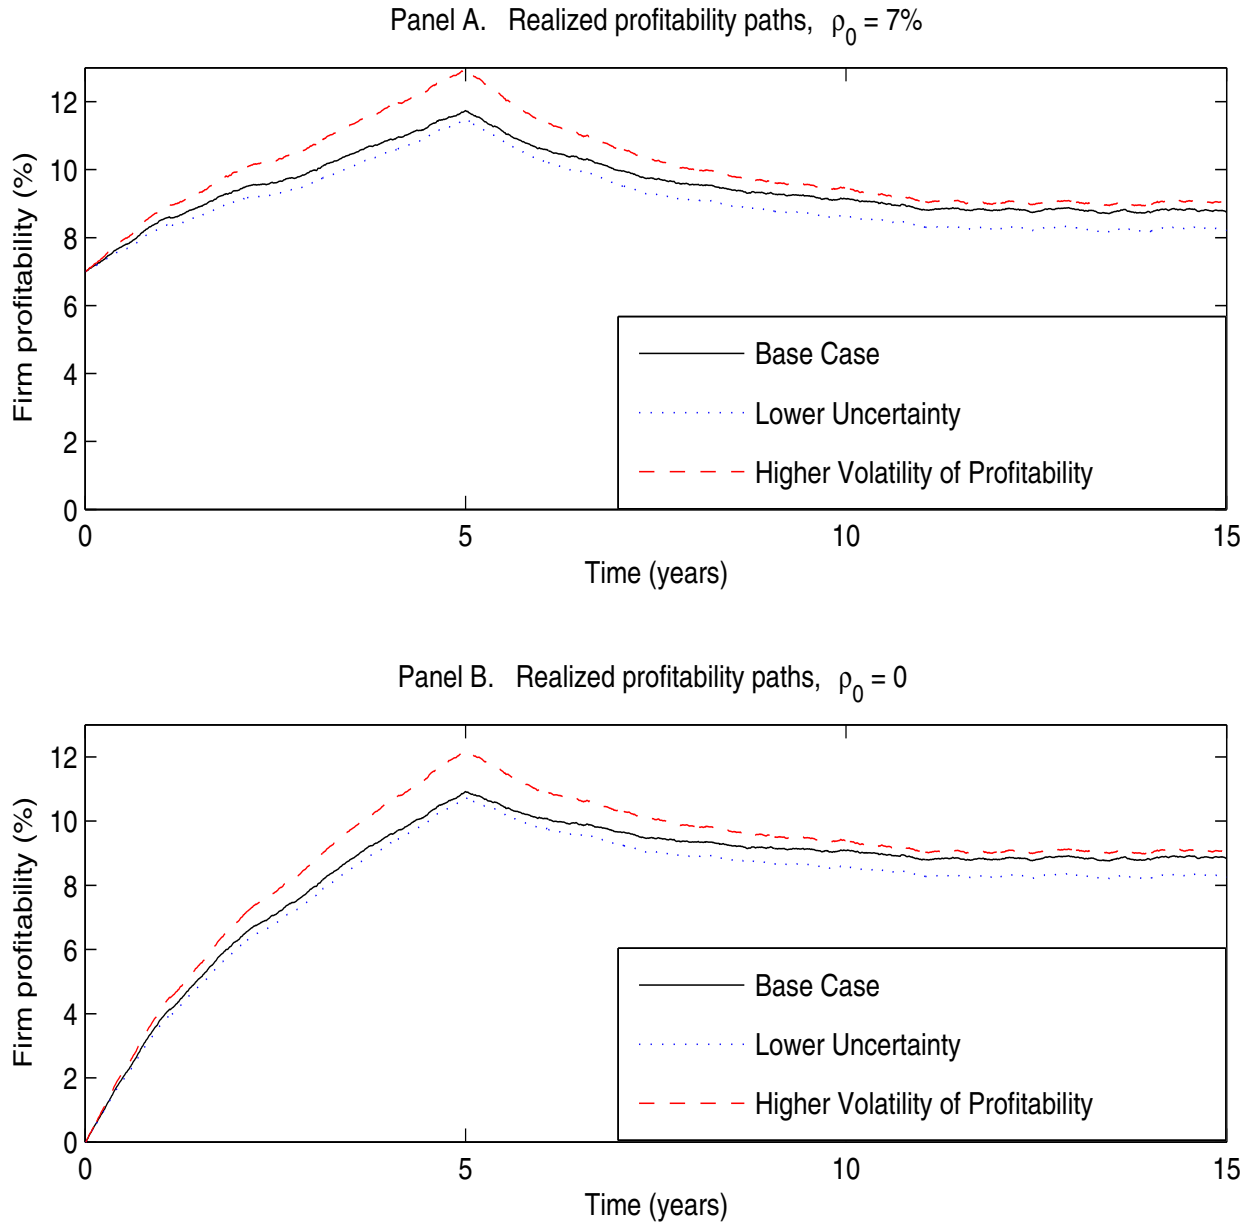


Figure 3. Model-Implied Realized Profitability Around an IPO. This figure plots the average paths of realized profitability, ρ_t , in percent per year, where the average is computed across all simulations of our model in which an IPO takes place at time $\tau = 5$. Given the large number of simulations, these average paths represent expected patterns in ρ_t conditional on an IPO. In Panel A, the initial profitability $\rho_0 = \hat{\rho}_0 = 7\%$; in Panel B, $\rho_0 = 0$. The solid line corresponds to the baseline case. The other two lines correspond to one-parameter deviations from the baseline case: uncertainty $\hat{\sigma}_0$ is reduced from 0.05 to 0.04 (dotted line), and volatility of profitability is increased from $(\sigma_{\rho,1}, \sigma_{\rho,2}) = (0.0584, 0.0596)$ to $\sigma_{\rho,1} = \sigma_{\rho,2} = 0.065$ (dashed line).

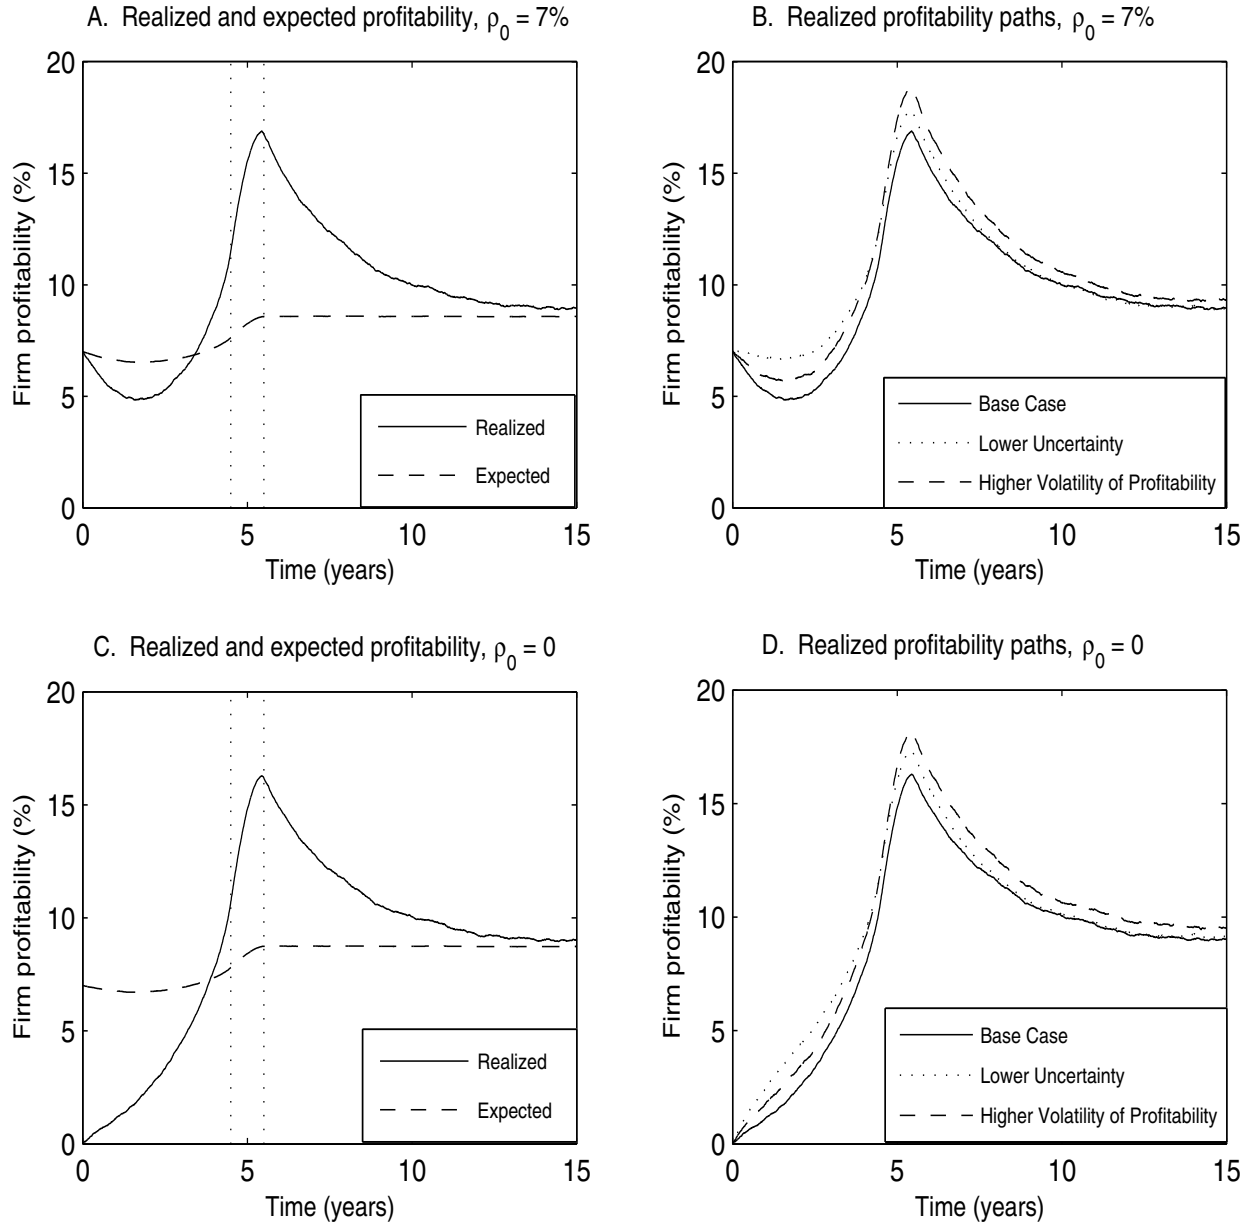


Figure 4. Model-Implied Profitability with Optimal IPO Time. This figure plots the average paths of realized profitability and expected average profitability, in percent per year, where the paths are averaged across all simulations of our model in which an IPO takes place at an optimally chosen time τ^* that falls within a one-year window centered at $\tau = 5$ years. Panels A and C are analogs of Panels A and B of Figure 2, whereas Panels B and D are analogs of Panels A and B of Figure 3. In Panels B and D, the solid line corresponds to the baseline case. The other two lines correspond to one-parameter deviations from the baseline case: uncertainty $\hat{\sigma}_0$ is reduced from 0.05 to 0.04 (dotted line), and volatility of profitability is increased from $(\sigma_{\rho,1}, \sigma_{\rho,2}) = (0.0584, 0.0596)$ to $\sigma_{\rho,1} = \sigma_{\rho,2} = 0.065$ (dashed line). In Panels A and B, the initial profitability $\rho_0 = \hat{\rho}_0 = 7\%$; in Panels C and D, $\rho_0 = 0$.

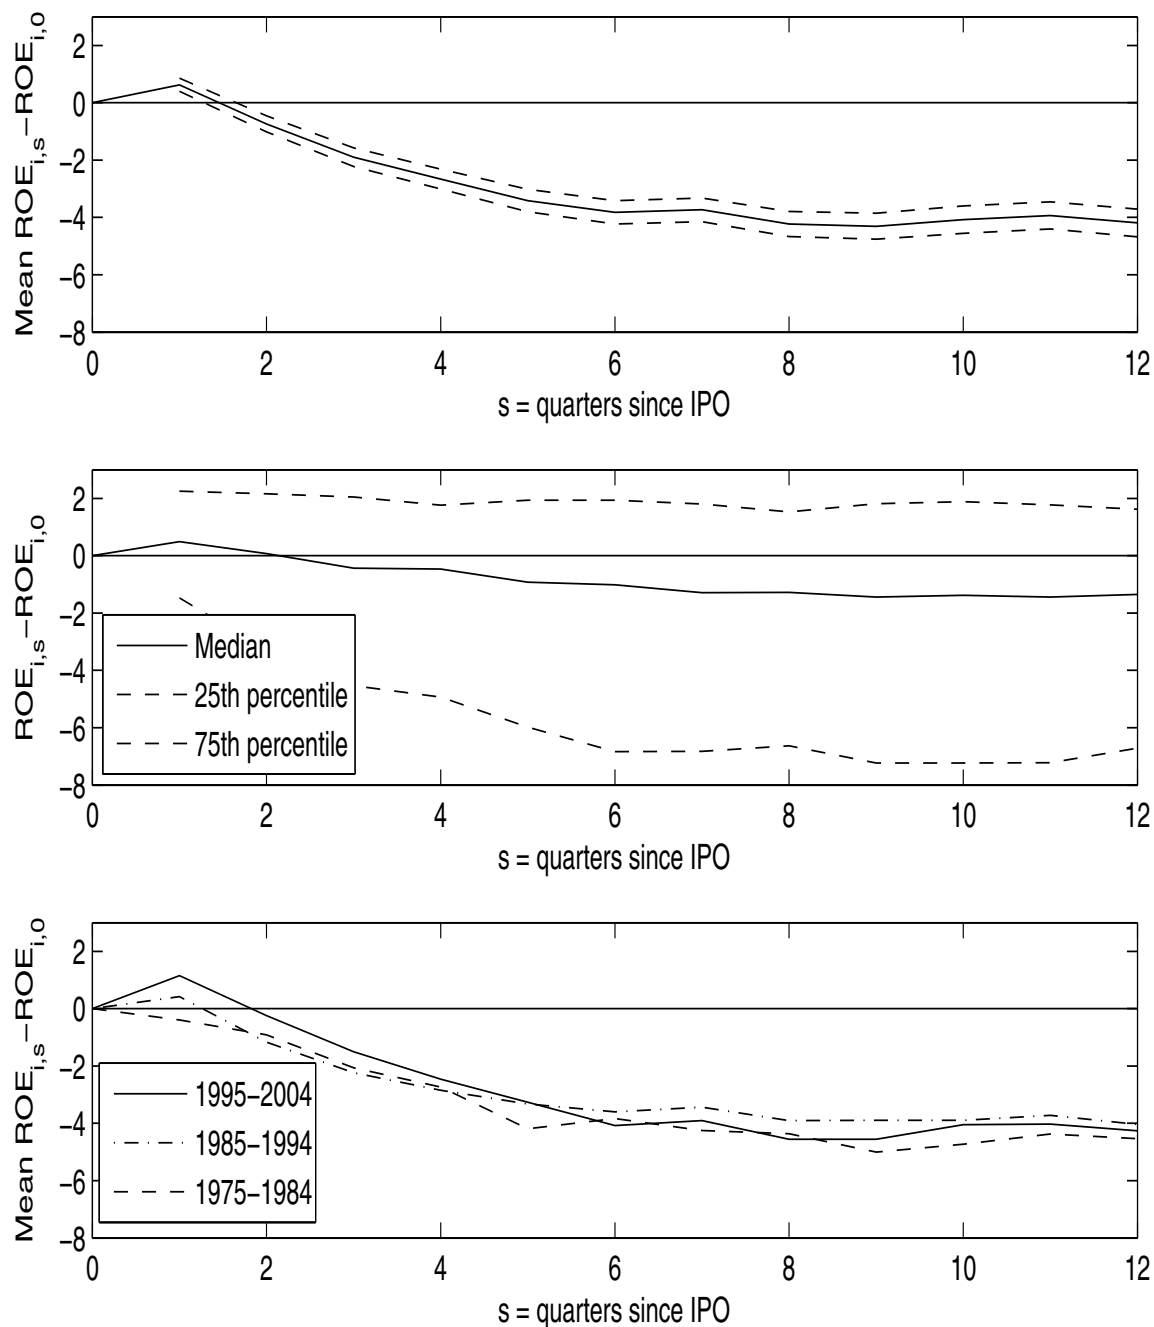


Figure 5. Post-IPO Changes in Profitability. This figure plots the post-IPO changes in firm profitability, measured as return on equity (ROE), for our sample of 7,183 IPOs in the U.S. from 1975–2004. Time 0 is the quarter of the IPO. $ROE_{i,s}$ is firm i 's profitability s quarters after its IPO, in percent per quarter. The top panel plots the equal-weighted average of $ROE_{i,s} - ROE_{i,0}$ across all firms for which both $ROE_{i,s}$ and $ROE_{i,0}$ can be computed (solid line), as well as the 95% confidence interval for the mean (dashed lines). The middle panel plots the median value of $ROE_{i,s} - ROE_{i,0}$ (solid line), as well as the 25th and 75th percentiles (dashed lines). The bottom panel plots the equal-weighted average of $ROE_{i,s} - ROE_{i,0}$ across IPOs in three sub-samples: 1975–1984, 1985–1994, and 1995–2004.

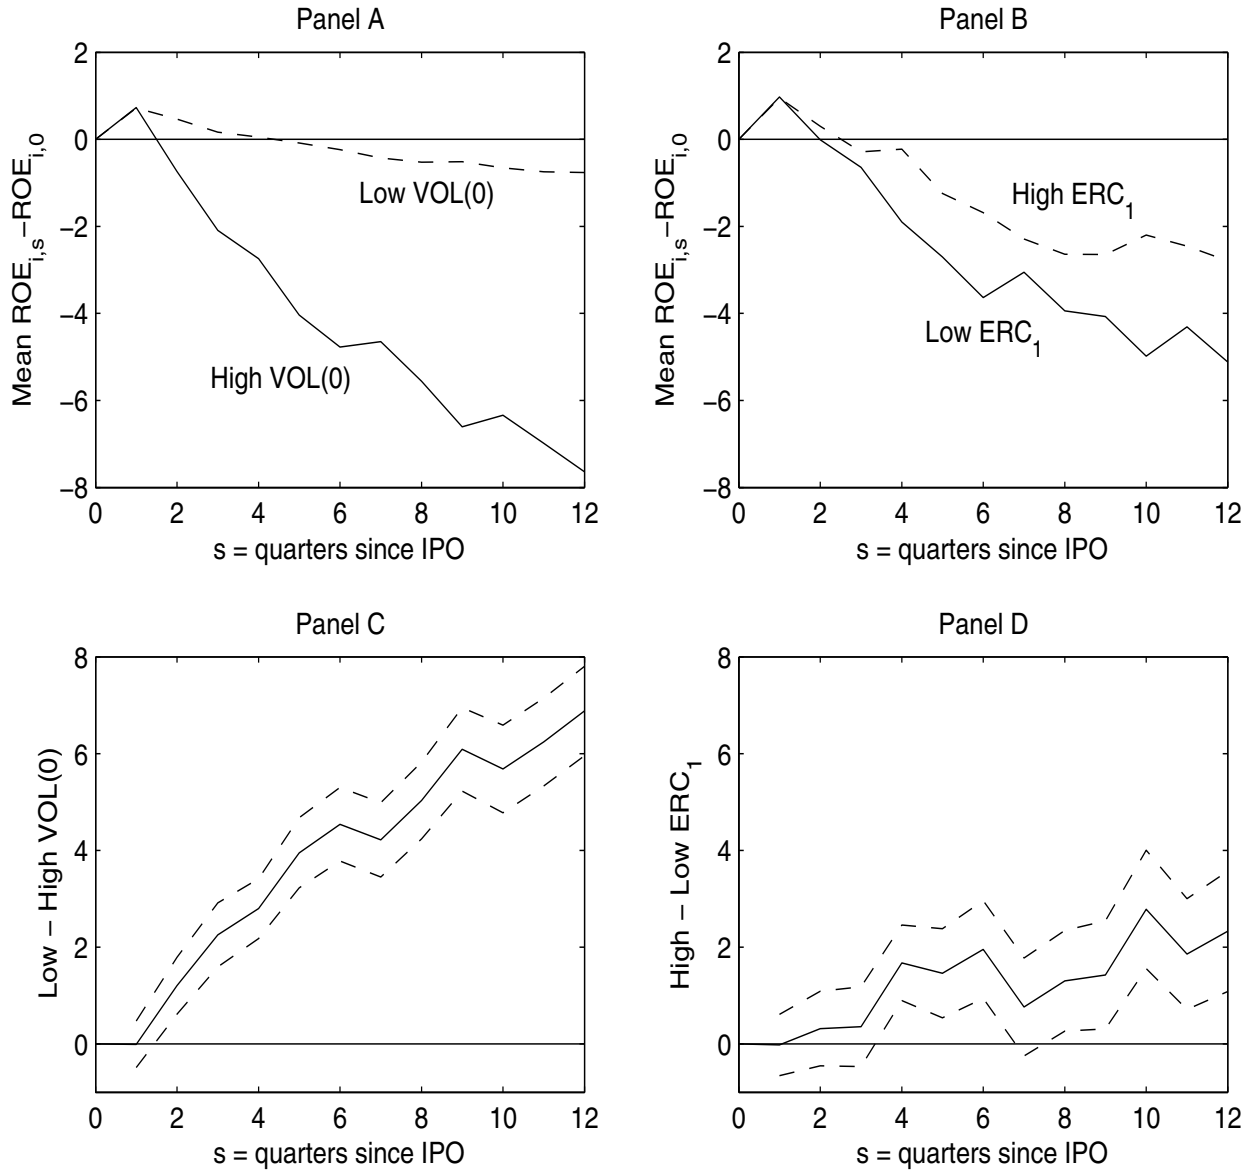


Figure 6. Post-IPO Changes in Profitability: Volatility vs. Uncertainty. We split our sample of 7,183 IPOs in 1975–2004 into high-volatility IPOs and low-volatility IPOs, and also into high- ERC_1 IPOs and low- ERC_1 IPOs. The left-hand panels split the sample using the median of $VOL(0)$, 5.31% per quarter. The right-hand panels split the sample using the median of ERC_1 , 2.21. ERC_1 measures firm i 's average stock price reaction to earnings surprises; $ROE_{i,s}$ is firm i 's profitability s quarters after its IPO, in percent per quarter; and $VOL(0)$ is the standard deviation of $ROE_{i,s}$ for $s = 0, \dots, 19$ quarters. Time 0 is the quarter of the IPO. Panels A and B plot the means of $ROE_{i,s} - ROE_{i,0}$ across the firms in the respective sub-samples split by volatility (Panel A) and ERC_1 (Panel B). Panel C plots the low volatility sub-sample's mean $ROE_{i,s} - ROE_{i,0}$ minus the high volatility sub-sample's mean $ROE_{i,s} - ROE_{i,0}$. Panel D plots the high ERC_1 sub-sample's mean $ROE_{i,s} - ROE_{i,0}$ minus the low ERC_1 sub-sample's mean $ROE_{i,s} - ROE_{i,0}$. The dashed lines denote the 95% confidence interval for this difference in differences.

Table 1
The Model-Implied Average Expected Post-IPO Drop in Profitability

Panel A shows the average expected post-IPO drop in profitability, computed at time 0 conditional on an IPO at time $\tau = 5$. Panel B shows the average volatility of the firm's stock returns, and Panel C reports the average expected excess return on the firm's stock. For any given combination of prior uncertainty, $\hat{\sigma}_0$, and the volatility of profitability, $\sigma_{\rho,1} = \sigma_{\rho,2}$, all three averages are computed across all admissible values of benefits of private control, α , and the prior mean, $\hat{\rho}_0$. The admissible values of α and $\hat{\rho}_0$ are subsets of the intervals [5%, 15%] and [-20%, 40%], respectively, that include only the sets of parameters for which it is optimal to start a private firm at time 0. The initial profitability is $\rho_0 = 0$ and all remaining parameters are at their baseline values.

		$\sigma_{\rho,1} = \sigma_{\rho,2}$ (% per year)									
		1	2	3	4	5	6	7	8	9	10
Panel A: Average Expected Drop in Profitability (% per year).											
	0	0.28	0.64	2.78	4.82	6.26	7.97	10.44	16.20	17.46	21.78
	1	-0.01	1.87	3.89	4.87	6.83	8.02	12.50	16.77	19.77	22.02
	2	-0.22	1.13	3.58	6.78	10.88	11.53	15.92	19.29	22.73	23.49
	3	-0.55	1.13	3.44	5.10	8.45	10.58	15.24	16.37	21.68	20.55
	4	-0.78	0.30	2.15	3.34	5.97	8.94	12.19	15.19	19.05	21.25
	5	-0.99	-0.12	1.17	2.78	4.61	7.12	9.67	13.86	14.29	18.94
(% p.a.)	6	-	-0.63	0.52	2.14	4.13	6.44	8.43	10.57	12.23	13.34
	7	-	-0.92	0.11	1.70	3.89	5.18	7.17	9.15	11.99	10.74
	8	-	-	-0.24	1.12	3.00	5.37	8.59	9.29	8.95	10.42
	9	-	-	-	-	2.27	4.12	6.76	-	-	-
	10	-	-	-	-	-	-	-	-	-	-
Panel B: Average Stock Return Volatility (% per year).											
	0	3.50	6.99	10.49	13.99	17.48	20.98	24.48	27.97	31.47	34.97
	1	4.83	7.92	11.16	14.50	17.90	21.33	24.78	28.24	31.70	35.18
	2	5.90	9.65	12.75	15.85	19.04	22.31	25.64	29.00	32.39	35.79
	3	6.30	10.97	14.48	17.59	20.66	23.77	26.95	30.19	33.47	36.78
	4	6.48	11.79	15.89	19.30	22.43	25.50	28.58	31.70	34.87	38.09
	5	6.57	12.29	16.94	20.77	24.13	27.27	30.34	33.40	36.50	39.62
	6	-	12.61	17.69	21.95	25.62	28.95	32.10	35.18	38.24	41.32
(% p.a.)	7	-	12.81	18.22	22.87	26.88	30.47	33.78	36.94	40.02	43.09
	8	-	-	18.62	23.58	27.93	31.79	35.31	38.60	41.77	44.86
	9	-	-	-	-	28.78	32.92	36.67	-	-	-
	10	-	-	-	-	-	-	-	-	-	-
Panel C: Average Expected Excess Stock Return (% per year).											
Any $\hat{\sigma}_0$		1.48	2.97	4.45	5.93	7.42	8.90	10.38	11.87	13.35	14.84

Table 2
Summary Statistics for the IPO Sample

Panel A contains summary statistics (means, standard deviations, percentiles) for the 7,183 firms in our sample of IPOs from 1975-2004. N is the number of firms for which the given variable can be calculated. t -stat is the t -statistic testing the hypothesis that the mean of the given variable is equal to zero. $ROE_{i,s}$ is the return on equity of firm i computed s quarters after the firm's IPO, in percent per quarter. $VOL(s_0)$ is the standard deviation of $ROE_{i,s}$ for $s = s_0, \dots, s_0 + 19$. ERC_1 is the average of the first 12 post-IPO stock price reactions to earnings surprises. ERC_1^+ is equal to ERC_1 when $ERC_1 > 0$ and missing otherwise. ERC_2 is the negative of the regression slope of earnings surprises on abnormal stock returns using firm i 's first 20 post-IPO quarters of earnings surprises. ERC_2^- is equal to ERC_2 when $ERC_2 < 0$ and missing otherwise. Panel B shows pairwise correlations computed across firms.

Panel A. Summary Statistics.								
Variable	N	Mean	Std. dev.	t -stat	Percentiles			
					25th	50th	75th	
$ROE_{i,0}$	6,149	-1.19	12.72	-7.3	-4.22	1.69	4.34	
$ROE_{i,12} - ROE_{i,0}$	4,254	-4.19	16.02	-17.1	-6.71	-1.35	1.62	
$VOL(0)$	4,744	8.19	7.70	73.2	2.52	5.31	11.44	
$VOL(13)$	2,843	7.89	8.19	51.4	2.27	4.63	10.58	
ERC_1	2,763	3.13	6.86	24.0	-1.08	2.21	6.82	
ERC_2	2,579	-0.035	0.067	-26.6	-0.064	-0.026	-0.002	
ERC_1^+	1,848	6.46	5.59	49.7	2.17	5.17	9.00	
ERC_2^-	2,001	-0.056	0.056	-44.7	-0.078	-0.040	-0.018	

Panel B. Cross-Sectional Correlations							
	$ROE_{i,12}$ $-ROE_{i,0}$	$VOL(0)$	$VOL(13)$	ERC_1	ERC_2	ERC_1^+	ERC_2^-
$ROE_{i,12} - ROE_{i,0}$	1.00						
$VOL(0)$	-0.33	1.00					
$VOL(13)$	-0.17	0.67	1.00				
ERC_1	0.07	-0.16	-0.10	1.00			
ERC_2	0.04	-0.13	-0.07	-0.06	1.00		
ERC_1^+	0.08	-0.25	-0.15	1.00	0.14	1.00	
ERC_2^-	0.15	-0.31	-0.19	0.16	1.00	0.30	1.00

Table 3
Cross-Sectional Regressions

This table reports OLS estimates of β from the model $ROE_{i,s} - ROE_{i,0} = \beta X_i + \epsilon_i$. The sample contains 7,183 IPO firms from 1975-2004 less any firms for which at least one variable is missing, for a total of N firms. $ROE_{i,s}$ is the return on equity of firm i computed s quarters after the firm's IPO, in percent per quarter. X_i contains combinations of the following variables: a constant, $VOL(s_0)$ (the standard deviation of $ROE_{i,s}$ for $s = s_0, \dots, s_0 + 19$), ERC_1 (the average of firm i 's first 12 post-IPO stock price reactions to earnings surprises), and ERC_2 (minus the regression slope of firm i 's earnings surprises on firm i 's abnormal stock returns around earnings announcements). The t -statistics, shown in parentheses, are computed by clustering the error terms in calendar time.

Panel A. One-Year Horizon. (Regressand: $ROE_{i,4} - ROE_{i,0}$)									
Constant	-2.67	0.71	0.44	-1.37	-0.41	0.52	0.65	0.38	0.62
	(-10.3)	(2.43)	(1.80)	(-4.29)	(-1.36)	(1.35)	(1.86)	(1.02)	(1.88)
$VOL(0)$		-0.26				-0.21	-0.17		
		(-12.2)				(-6.94)	(-5.41)		
$VOL(5)$			-0.19					-0.17	-0.13
			(-9.40)					(-5.66)	(-4.34)
ERC_1				0.10		0.06		0.04	
				(3.41)		(1.98)		(1.17)	
ERC_2					5.23		2.75		6.97
					(1.82)		(0.94)		(2.29)
R^2	0.000	0.033	0.024	0.004	0.001	0.026	0.014	0.019	0.013
N	5,777	4,399	3,628	2,649	2,474	2,338	2,425	1,947	2,112
Panel B. Three-Year Horizon. (Regressand: $ROE_{i,12} - ROE_{i,0}$)									
Constant	-4.19	1.32	-0.52	-4.42	-2.70	1.91	1.32	-0.45	0.04
	(-13.7)	(3.74)	(-1.52)	(-9.62)	(-6.76)	(3.28)	(2.91)	(-0.87)	(0.07)
$VOL(0)$		-0.70				-0.84	-0.66		
		(-22.4)				(-18.9)	(-15.3)		
$VOL(13)$			-0.27					-0.34	-0.29
			(-8.68)					(-8.24)	(-6.86)
ERC_1				0.15		0.02		0.02	
				(3.23)		(0.45)		(0.45)	
ERC_2					9.14		1.81		12.14
					(2.13)		(0.44)		(2.66)
R^2	0.000	0.106	0.029	0.005	0.002	0.147	0.093	0.049	0.036
N	4,254	4,229	2,541	2,259	2,368	2,255	2,368	1,365	1,517

Table 4
Cross-Sectional Regressions, Excluding ERCs with Unpredicted Signs

This table reports OLS estimates of β from the model $ROE_{i,s} - ROE_{i,0} = \beta_s X_i + \epsilon_i$. The sample contains 7,183 IPO firms from 1975-2004 less any firms for which at least one variable is missing, for a total of N firms. $ROE_{i,s}$ is the return on equity of firm i computed s quarters after the firm's IPO, in percent per quarter. X_i contains combinations of the following variables: a constant, $VOL(s_0)$ (the standard deviation of $ROE_{i,s}$ for $s = s_0, \dots, s_0 + 19$), ERC_1^+ (the average of firm i 's first 12 post-IPO stock price reactions to earnings surprises, excluding negative values), and ERC_2^- (minus the regression slope of firm i 's earnings surprises on firm i 's abnormal stock returns around earnings announcements, excluding positive values). The t -statistics, shown in parentheses, are computed by clustering the error terms in calendar time.

Panel A. One-Year Horizon. (Regressand: $ROE_{i,4} - ROE_{i,0}$)									
Constant	-2.67 (-10.3)	0.71 (2.43)	0.44 (1.80)	-1.95 (-5.38)	-0.06 (-0.16)	-0.40 (-0.78)	0.80 (1.95)	-0.74 (-1.52)	0.65 (1.62)
$VOL(0)$		-0.26 (-12.2)				-0.13 (-3.68)	-0.18 (-4.88)		
$VOL(5)$			-0.19 (-9.40)					-0.10 (-2.92)	-0.09 (-2.54)
ERC_1^+				0.18 (4.34)		0.11 (2.62)		0.12 (2.74)	
ERC_2^-					11.21 (2.90)		4.78 (1.16)		12.90 (3.13)
R^2	0.000	0.033	0.024	0.011	0.004	0.017	0.018	0.015	0.014
N	5,777	4,399	3,628	1,774	1,930	1,582	1,893	1,323	1,664
Panel B. Three-Year Horizon. (Regressand: $ROE_{i,12} - ROE_{i,0}$)									
Constant	-4.19 (-13.7)	1.32 (3.74)	-0.52 (-1.52)	-4.78 (-7.83)	-0.81 (-1.58)	1.47 (1.76)	1.55 (3.21)	-0.47 (-0.75)	0.61 (1.04)
$VOL(0)$		-0.70 (-22.4)				-0.71 (-13.3)	-0.54 (-10.8)		
$VOL(13)$			-0.27 (-8.68)					-0.29 (-6.55)	-0.22 (-4.88)
ERC_1^+				0.22 (3.32)		0.00 (0.00)		0.01 (0.19)	
ERC_2^-					36.59 (6.67)		18.89 (3.40)		28.85 (5.00)
R^2	0.000	0.106	0.029	0.007	0.024	0.116	0.081	0.045	0.047
N	4,254	4,229	2,541	1,532	1,853	1,530	1,853	929	1,206

Table 5**Counterpart of Table 4, Controlling for Leverage, Size, and Industry Fixed Effects**

This table reports OLS estimates of β from the model $ROE_{i,s} - ROE_{i,0} = \beta_s X_i + \gamma_s$ controls $i + \epsilon_i$. The only difference from Table 4 is that we include three controls: (i) leverage (LEV), the firm's debt divided by total assets at the end of the IPO quarter, (ii) size (SIZE), the log ratio of the firm's assets to total assets of firms in Compustat at the end of the IPO's fiscal year, and (iii) industry fixed effects, defined at the 2-digit SIC code level.

Panel A. One-Year Horizon. (Regressand: $ROE_{i,4} - ROE_{i,0}$)								
Constant	-13.29 (-1.27)	-13.34 (-1.35)	6.07 (1.74)	-9.17 (-0.93)	4.84 (1.38)	-11.89 (-1.21)	4.66 (1.31)	-9.06 (-0.96)
$VOL(0)$	-0.22 (-9.79)				-0.11 (-2.99)	-0.17 (-4.46)		
$VOL(5)$		-0.17 (-7.57)					-0.08 (-2.35)	-0.08 (-2.17)
ERC_1^+			0.16 (3.69)		0.11 (2.39)		0.11 (2.49)	
ERC_2^-				9.02 (2.28)		2.91 (0.69)		10.29 (2.42)
LEV	0.54 (0.62)	-0.52 (-0.59)	-2.25 (-1.82)	-1.17 (-0.98)	-1.43 (-1.12)	-0.54 (-0.44)	-2.79 (-2.15)	-2.26 (-1.86)
$SIZE$	0.26 (1.83)	0.28 (1.94)	0.65 (2.80)	0.33 (1.49)	0.50 (2.12)	0.15 (0.69)	0.43 (1.79)	0.27 (1.23)
R^2	0.061	0.045	0.061	0.045	0.070	0.056	0.080	0.058
N	4,175	3,440	1,757	1,911	1,565	1,874	1,307	1,645
Panel B. Three-Year Horizon. (Regressand: $ROE_{i,12} - ROE_{i,0}$)								
Constant	-2.78 (-0.18)	-0.32 (-0.02)	18.08 (3.39)	2.25 (0.17)	10.91 (2.13)	-3.35 (-0.26)	9.21 (1.80)	9.69 (2.31)
$VOL(0)$	-0.72 (-20.91)				-0.73 (-13.38)	-0.54 (-10.46)		
$VOL(13)$		-0.27 (-8.13)					-0.30 (-6.63)	-0.21 (-4.67)
ERC_1^+			0.14 (2.14)		-0.05 (-0.85)		-0.03 (-0.55)	
ERC_2^-				34.89 (6.25)		17.65 (3.12)		28.26 (4.88)
LEV	2.00 (1.58)	0.19 (0.14)	-5.98 (-3.12)	-2.02 (-1.24)	-2.41 (-1.32)	-0.86 (-0.54)	-5.01 (-2.75)	-3.94 (-2.45)
$SIZE$	-0.21 (-1.00)	0.05 (0.20)	1.56 (4.38)	0.67 (2.27)	0.69 (2.01)	0.29 (0.99)	0.65 (1.89)	0.61 (1.99)
R^2	0.124	0.068	0.073	0.077	0.172	0.128	0.117	0.053
N	4,009	2,390	1,514	1,833	1,512	1,833	913	1,188

Appendix.

Profitability, ROE_{is} , equals $[I_{is} + DT_{is}] / BE_{is}$. The subscript s denotes the s -th fiscal quarter after the fiscal quarter of firm i 's IPO. The fiscal quarter containing the IPO is quarter zero. I_{is} equals the income before extraordinary items available for common stock (Compustat quarterly item 25) for firm i in quarter s . DT_{is} equals deferred taxes from income account (Compustat quarterly item 35); we impute a zero value if this item is missing. BE_{is} is the book value of equity of firm i in quarter s . BE_{is} is calculated either from the previous fiscal quarter, previous fiscal year, current fiscal quarter, or current fiscal year, taken in that order depending on availability. To mitigate the backfilling bias, we do not use pre-IPO Compustat data. Following Fama and French (1993), book value of equity equals stockholders' equity plus deferred taxes minus book value of preferred stock. If any of these three items is missing, then book value of equity is treated as missing. We treat negative or zero values of BE as missing. Stockholders' equity equals either "total stockholders' equity" (quarterly item 60, annual item 216), "total common equity" (quarterly item 59, annual item 60) + "carrying value of preferred stock" (quarterly item 55, annual item 130), "total assets" (quarterly item 44, annual item 6) - "total liabilities" (quarterly item 54, annual item 181), or missing, in that order depending on availability. Deferred taxes equals "deferred tax and investment tax credit" (quarterly item 52, annual item 35), or if that is missing, then zero. Annual book value of preferred stock equals either "redemption value of preferred stock" (annual item 56), "liquidating value of preferred stock" (annual item 10), "carrying value of preferred stock" (annual 130), or zero, in that order depending on availability. Quarterly book value of preferred stock equals "book value of preferred stock" (quarterly item 55), or zero if item 55 is missing. We eliminate firm-quarter observations where ROE_{is} is outside $[-100\%, +100\%]$.

Abnormal stock return, AR_{it} , is the cumulative return of stock i in excess of stock i 's industry, starting one day before the stock's t -th post-IPO earnings announcement and ending one day after the same announcement. Since the industry portfolios were constructed using Compustat SIC codes, we link firms to industries using the most recent annual Compustat SIC code (item 324), soonest future Compustat annual SIC code, most recent CRSP SIC code (SICCD), or soonest future CRSP SIC code, in that order depending on availability. Earnings announcement date is variable REPDATS from the IBES unadjusted actuals file.

Earnings per share, EPS_{it} , is the quarterly EPS of firm i announced in its t -th post-IPO earnings announcement (variable VALUE in the IBES unadjusted actuals file). $E[EPS_{it}]$ is the mean of all analyst forecasts of EPS_{it} using IBES's last pre-announcement set of forecasts for the given fiscal quarter (variable MEANEST in the IBES unadjusted summary file). We eliminate observations for which the earnings announcement date is more than 60 days after the most recent set of earnings forecasts (roughly 1% of observations are eliminated).

Theoretical Results.

This section of the Appendix contains the formulas that we refer to in the text. The proofs of all propositions along with some additional technical material are contained in the Technical Appendix, which is available on the authors' websites.

Learning: The agents' posterior beliefs about $\bar{\rho}$ are given by the normal distribution in equation (15), whose posterior mean and variance evolve over time according to

$$d\hat{\rho}_t = \hat{\sigma}_t^2 \frac{\phi}{\sigma_{\rho,2}} d\widehat{X}_{2,t}; \quad \hat{\sigma}_t^2 = \frac{1}{\frac{1}{\sigma_0^2} + \left(\frac{\phi}{\sigma_{\rho,2}}\right)^2 t}, \quad (27)$$

and $d\widehat{X}_{2,t}$ is a Brownian motion given by the normalized expectation error.

Market value: Let $\sigma_\pi = (\sigma_{\pi,1}, 0)$ and $\sigma_\rho = (\sigma_{\rho,1}, \sigma_{\rho,2})$. The firm's market value is

$$M_t = B_t e^{Q_0(T-t) + Q_1(T-t)\rho_t + Q_2(T-t)\hat{\rho}_t + \frac{1}{2}Q_2(T-t)^2\hat{\sigma}_t^2}, \quad (28)$$

where

$$\begin{aligned} Q_0(s) &= -rs + \frac{\sigma_\rho \sigma'_\rho}{2\phi^2} Q_3(s) - \frac{\sigma_\pi \sigma'_\rho}{\phi} Q_2(s); & Q_1(s) &= \frac{1}{\phi} (1 - e^{-\phi s}) > 0; \\ Q_2(s) &= s - Q_1(s) > 0; & Q_3(s) &= s + \frac{1 - e^{-2\phi s}}{2\phi} - 2Q_1(s). \end{aligned}$$

Selling the whole firm in an IPO is optimal: We provide two different proofs.

Proof #1. Since the benefits of private control disappear after the IPO at time τ , the entrepreneur's post-IPO consumption is financed only from asset returns. Given the state price density in equation (11), the entrepreneur's optimal consumption at any post-IPO time u , $\tau < u < T$, follows from the results in Cox and Huang (1989):

$$c_u = e^{-\frac{\beta}{\gamma}(u-\tau)} \left(\frac{\pi_u}{\pi_\tau}\right)^{-\frac{1}{\gamma}} \lambda^{-\frac{1}{\gamma}}, \quad W_T = e^{-\beta(T-\tau)} \left(\frac{\pi_T}{\pi_\tau}\right)^{\frac{1}{\gamma}} \lambda^{-\frac{1}{\gamma}}.$$

Hence, the optimal consumption is driven only by $dX_{1,t}$ shocks, which are perfectly correlated with the returns on publicly-traded stocks (the stock market). Investing any amount in the entrepreneur's (tiny) firm would make the entrepreneur's consumption driven also by the firm's idiosyncratic shocks $dX_{2,t}$, which would make the consumption path suboptimal.

Proof #2. We solve the entrepreneur's portfolio choice problem at time τ and show that the optimal retention ratio θ_τ^f is zero. At time τ , the entrepreneur sells the firm for M_τ in (28) and buys back the optimal fraction θ_τ^f of the same firm. His remaining wealth is optimally invested in stocks (fraction θ_τ) and bonds (fraction $1 - \theta_\tau - \theta_\tau^f$). The entrepreneur solves

$$V(W_\tau, \tau) = \max_{\{c_t, \theta_t, \theta_t^f\}_{t=\tau}^{t=T}} E_\tau \left[\int_\tau^T e^{-\beta(t-\tau)} \frac{c_t^{1-\gamma}}{1-\gamma} dt + \eta e^{-\beta(T-t)} \frac{W_T^{1-\gamma}}{1-\gamma} \right], \quad (29)$$

subject to the budget constraint $dW_t = W_t (\theta_t dR_t + \theta_t^f \frac{dM}{M_t} + (1 - \theta_t - \theta_t^f) r dt) - c_t dt$. As in proof #1, c_t is financed only from asset returns. The return on stocks follows the process

$$dR_t = (r + \mu_R) dt + \sigma_R dX_{1,t},$$

where the risk premium $\mu_R = \sigma_\pi \sigma_R$. The process for the firm's post-IPO stock return ($\frac{dM}{M_t}$) is obtained by applying Ito's Lemma to M_t (see Proposition 2 in Pástor and Veronesi, 2003). Solving the Hamilton, Jacobi, Bellman equation implied by equation (29), we find

$$\theta_t^f = 0 \quad \text{and} \quad \theta_t = \frac{\mu_R}{\gamma \sigma_R^2}. \quad (30)$$

Utility from selling the firm: The solution to (29) is the value function

$$V(W_\tau, \tau) = \frac{W_\tau^{1-\gamma}}{1-\gamma} \left(\frac{\left(1 + \eta \frac{1-\gamma}{\gamma} \left(r - \frac{\beta}{1-\gamma} + \frac{1}{2} \frac{1}{\gamma} \sigma_{\pi,1}^2\right)\right) e^{\frac{1-\gamma}{\gamma} \left(r - \frac{\beta}{1-\gamma} + \frac{1}{2} \frac{1}{\gamma} \sigma_{\pi,1}^2\right) (T-\tau)} - 1}{\frac{1-\gamma}{\gamma} \left(r - \frac{\beta}{1-\gamma} + \frac{1}{2} \frac{1}{\gamma} \sigma_{\pi,1}^2\right)} \right)^\gamma. \quad (31)$$

Utility from keeping the firm: The utility from owning the firm from τ to T is given by

$$V^O(B_\tau, \tau) = \frac{B_\tau^{1-\gamma}}{1-\gamma} \left\{ \alpha^{1-\gamma} \int_\tau^T Z^O(\rho_\tau, \hat{\rho}_\tau, \hat{\sigma}_\tau^2; u - \tau) du + \eta Z^O(\rho_\tau, \hat{\rho}_\tau, \hat{\sigma}_\tau^2; T - \tau) \right\}, \quad (32)$$

where

$$Z^O(\rho_t, \hat{\rho}_t, \hat{\sigma}_t^2; s) = e^{\bar{Q}_0(s) + (1-\gamma)Q_1(s)\rho_t + (1-\gamma)Q_2(s)\hat{\rho}_t + \frac{1}{2}(1-\gamma)^2 Q_3(s)^2 \hat{\sigma}_t^2} \quad (33)$$

in which $Q_i(\cdot)$ are given above and $\bar{Q}_0(s) = -\beta s + (1-\gamma)^2 \frac{\sigma_\rho \sigma'_\rho}{2\phi^2} Q_3(s)$.

IPO condition: An IPO takes place at time τ if and only if

$$f(T - \tau, \hat{\sigma}_\tau, \sigma_\rho) \leq \alpha^{1-\gamma} \int_\tau^T \hat{Z}(\rho_\tau, \hat{\rho}_\tau, \hat{\sigma}_\tau, \sigma_\rho; u - \tau; T) du, \quad (34)$$

where

$$f(T - \tau, \hat{\sigma}_\tau, \sigma_\rho) = e^{(1-\gamma) \left(-\left(r - \frac{\beta}{1-\gamma}\right) (T-\tau) + \gamma \frac{\sigma_\rho \sigma'_\rho}{2\phi^2} Q_3(T-\tau) - \frac{\sigma_\pi \sigma'_\rho}{\phi} Q_2(T-\tau) \right) + \frac{1}{2} \gamma (1-\gamma) Q_2(T-\tau)^2 \hat{\sigma}_\tau^2} g(T - \tau) - \eta \hat{Z}(\rho_\tau, \hat{\rho}_\tau, \hat{\sigma}_\tau; u - \tau; T) = e^{\hat{Q}_0(u-\tau; T) + (1-\gamma)\hat{Q}_1(u-\tau; T)\rho_\tau + (1-\gamma)\hat{Q}_2(u-\tau; T)\hat{\rho}_\tau + \frac{1}{2}(1-\gamma)^2 \hat{Q}_3(u-\tau; T)\hat{\sigma}_\tau^2}$$

Note that f decreases in both $\hat{\sigma}_\tau$ and $\sigma_{\rho,2}$, \hat{Z} increases in both ρ_τ and $\hat{\rho}_\tau$, and $\hat{Z} > 0$. Above,

$$g(T - t) = \left(\frac{\left(1 + \eta \frac{1-\gamma}{\gamma} \left(r - \frac{\beta}{1-\gamma} + \frac{1}{2} \frac{1}{\gamma} \sigma_{\pi,1}^2\right)\right) e^{\frac{1-\gamma}{\gamma} \left(r - \frac{\beta}{1-\gamma} + \frac{1}{2} \frac{1}{\gamma} \sigma_{\pi,1}^2\right) (T-t)} - 1}{\frac{1-\gamma}{\gamma} \left(r - \frac{\beta}{1-\gamma} + \frac{1}{2} \frac{1}{\gamma} \sigma_{\pi,1}^2\right)} \right)^\gamma$$

and

$$\begin{aligned} \hat{Q}_0(u - \tau; T) &= \bar{Q}_0(u - \tau) - \bar{Q}_0(T - \tau) \\ \hat{Q}_1(u - \tau; T) &= Q_1(u - \tau) - Q_1(T - \tau) < 0 \\ \hat{Q}_2(u - \tau; T) &= Q_2(u - \tau) - Q_2(T - \tau) < 0 \\ \hat{Q}_3(u - \tau; T) &= Q_2(u - \tau)^2 - Q_2(T - \tau)^2 < 0 \end{aligned}$$

Equation (34) is equivalent to the IPO condition in equation (16).

Endogenous IPO Cutoff Rule: The IPO condition (34) can be restated as follows:

$$f(T - \tau, \hat{\sigma}_\tau, \sigma_\rho) \leq h(x_\tau, \hat{\rho}_\tau) \equiv \alpha^{1-\gamma} \int_\tau^T \bar{Z}(x_\tau, \hat{\rho}_\tau, \hat{\sigma}_\tau^2, u - \tau, T) du, \quad (35)$$

where $\bar{Z}(x_\tau, \hat{\rho}_\tau, \hat{\sigma}_\tau, u - \tau, T) = e^{\hat{Q}_0(u-\tau, T) + (1-\gamma)(\hat{Q}_1(u-\tau, T)x_\tau + (u-T)\hat{\rho}_\tau) + \frac{1}{2}(1-\gamma)^2 \hat{Q}_3(u-\tau, T)\hat{\sigma}_\tau^2}$. The function $h(x_\tau, \hat{\rho}_\tau)$ is monotonically increasing in x_τ and $\hat{\rho}_\tau$. Assuming that $f(T - \tau, \hat{\sigma}_\tau, \sigma_\rho)$ is sufficiently large, we can define the cutoff $\underline{\rho}(x_\tau; \hat{\sigma}_\tau, \sigma_\rho)$ such that

$$h(x_\tau, \underline{\rho}(x_\tau; \hat{\sigma}_\tau, \sigma_\rho)) = f(T - \tau, \hat{\sigma}_\tau, \sigma_\rho). \quad (36)$$

If $f(T - \tau, \hat{\sigma}_\tau, \sigma_\rho)$ is too low for such a cutoff to exist, we set $\underline{\rho}(x_\tau; \hat{\sigma}_\tau, \sigma_\rho) = -\infty$.

Expected drop in profitability: This expectation is given in equation (18), where

$$k(x_\tau, \tau; t, x_t, \hat{\rho}_t, \hat{\sigma}_t^2) = \frac{\underline{\rho}(x_\tau) - \hat{\rho}_t - a(t, \tau; \hat{\sigma}_t^2)(x_\tau - e^{-\phi(\tau-t)}x_t)}{\sqrt{(\hat{\sigma}_t^2 - \hat{\sigma}_\tau^2)(1 - b(t, \tau; \hat{\sigma}_t^2)^2)}} \quad (37)$$

$$\mu_x = e^{-\phi(\tau-t)}x_t; \quad \sigma_x^2 = \frac{1 - e^{2\phi(\tau-t)}}{2\phi}(\sigma_{\rho,1}^2 + \sigma_{\rho,2}^2) + (e^{-2\phi(\tau-t)}\hat{\sigma}_t^2 - \hat{\sigma}_\tau^2) \quad (38)$$

and $a(t, \tau; \hat{\sigma}_t^2)$ and $b(t, \tau; \hat{\sigma}_t^2)$ are given by

$$a(t, \tau; \hat{\sigma}_t^2) = \frac{\hat{\sigma}_\tau^2 - e^{-\phi(\tau-t)}\hat{\sigma}_t^2}{\frac{1 - e^{2\phi(\tau-t)}}{2\phi}(\sigma_{\rho,1}^2 + \sigma_{\rho,2}^2) + (e^{-2\phi(\tau-t)}\hat{\sigma}_t^2 - \hat{\sigma}_\tau^2)}$$

$$b(t, \tau; \hat{\sigma}_t^2) = \frac{\hat{\sigma}_\tau^2 - e^{-\phi(\tau-t)}\hat{\sigma}_t^2}{\sqrt{\frac{1 - e^{2\phi(\tau-t)}}{2\phi}(\sigma_{\rho,1}^2 + \sigma_{\rho,2}^2) + (e^{-2\phi(\tau-t)}\hat{\sigma}_t^2 - \hat{\sigma}_\tau^2)} \sqrt{\hat{\sigma}_t^2 - \hat{\sigma}_\tau^2}$$

The Decision to Start a Private Firm: At time $t = 0$, the entrepreneur is endowed with a patent-protected technology and the initial wealth W_0 . To produce a stream of profits, the technology requires an initial investment of $B_0 = W_0$. The entrepreneur has three choices:

- (A) Start a private firm. (Invest W_0 in the technology to start production, keep the firm.)
- (B) Sell the patent to investors. (Invest W_0 in the technology to start production, sell it to investors for its fair market value M_0 , invest M_0 in stocks and bonds.)
- (C) Discard the patent. (Invest W_0 in stocks and bonds.)

The entrepreneur makes a utility-maximizing choice between (A), (B), and (C). Under choice (C), his expected utility is $V(B_0, 0)$. Under choice (B), his utility is $V(M_0, 0)$. Under choice (A), his expected utility, which we denote by $V_0^O(B_0, 0)$, is given by

$$V_0^O(B_0, 0) = E_0 \left[\int_0^\tau e^{-\beta t} \frac{(\alpha B_t)^{1-\gamma}}{1-\gamma} dt \right] + e^{-\beta\tau} E_0 \left[V(M_\tau, \tau) | \hat{\rho}_\tau \geq \underline{\rho} \right] \Pr(\hat{\rho}_\tau \geq \underline{\rho}) \quad (39)$$

$$+ e^{-\beta\tau} E_0 \left[V^O(B_\tau, \tau) | \hat{\rho}_\tau < \underline{\rho} \right] \Pr(\hat{\rho}_\tau < \underline{\rho}),$$

where ‘‘Pr’’ stands for ‘‘probability’’ as of time 0. The first term on the right-hand side reflects the benefits of private control consumed before time τ . The second term is the present value of expected utility conditional on an IPO taking place at time τ . The third term is the utility obtained if no IPO takes place. The closed-form solution for $V_0^O(B_0, 0)$ is given in the Technical Appendix. (A) is the optimal choice if and only if

$$V_0^O(B_0, 0) > \max \{V(M_0, 0), V(B_0, 0)\}. \quad (40)$$

REFERENCES

- Asquith, Paul, and David W. Mullins, 1986, "Equity issues and offering dilution", *Journal of Financial Economics* 15, 61–89.
- Ball, Ray, and Lakshmanan Shivakumar, 2006, "Earnings quality at initial public offerings", Working paper, University of Chicago.
- Barry, Christopher B., Chris J. Muscarella, John W. Peavy III, and Michael R. Vetsuypens, 1990, "The role of venture capital in the creation of public companies," *Journal of Financial Economics* 27, 447–471.
- Beaver, William H., 1970, "The time series behavior of earnings", *Journal of Accounting Research* 8, 62–89.
- Benninga, Simon, Mark Helmantel, and Oded Sarig, 2005, "The timing of initial public offerings", *Journal of Financial Economics* 75, 115–132.
- Bharath, Sreedhar T., and Amy K. Dittmar, 2006, "To be or not to be (public)", Working paper, University of Michigan.
- Black, Bernard S., and Ronald J. Gilson, 1998, "Venture capital and the structure of capital markets: Banks versus stock markets", *Journal of Financial Economics* 47, 243–277.
- Bodnaruk, Andriy, Eugene Kandel, Massimo Massa, and Andrei Simonov, 2006, "Shareholder diversification and the decision to go public", *Review of Financial Studies*, forthcoming.
- Brau, James C., and Stanley E. Fawcett, 2006, "Initial public offerings: An analysis of theory and practice", *Journal of Finance* 61, 399–436.
- Brav, Alon, Christopher C. Géczy, and Paul A. Gompers, 2000, "Is the abnormal return following equity issuances anomalous?", *Journal of Financial Economics* 56, 209–249.
- Brav, Alon, and Paul A. Gompers, 1997, "Myth or reality? The long-run underperformance of initial public offerings: Evidence from venture and nonventure capital-backed companies", *Journal of Finance* 52, 1791–1821.
- Chan, Louis K.C., Narasimhan Jegadeesh, and Josef Lakonishok, 1995, Evaluating the performance of value versus glamour stocks: The impact of selection bias, *Journal of Financial Economics* 38, 269–296.
- Chemmanur, Thomas J., 1993, "The pricing of initial public offerings: A dynamic model with information production," *Journal of Finance* 48, 285–304.
- Chemmanur, Thomas J., and Paolo Fulghieri, 1999, "A theory of the going-public decision," *Review of Financial Studies* 12, 249–279.
- Chemmanur, Thomas J., Shan He, and Debarshi Nandy, 2006, "The going-public decision and the product market: Empirical evidence," Working paper, Boston College.
- Clementi, Gian Luca, 2002, "IPOs and the growth of firms," Working paper, NYU.

- Cox, John C., and Chi-fu Huang, 1989, "Optimal consumption and portfolio choices when asset prices follow a diffusion process", *Journal of Economic Theory* 49, 33–83.
- DeGeorge, Francois, and Richard Zeckhauser, 1993, "The reverse LBO decision and firm performance: Theory and evidence", *Journal of Finance* 48, 1323–1348.
- Dyck, Alexander, and Luigi Zingales, 2004, "Private benefits of control: An international comparison", *Journal of Finance* 59, 537–600.
- Easton, Peter D., and Mark E. Zmijewski, 1989, "Cross-sectional variation in the stock market response to the announcement of accounting earnings," *Journal of Accounting and Economics* 11, 117–141.
- Fama, Eugene F. and Kenneth R. French, 1993, "Common risk factors in the returns on stocks and bonds", *Journal of Financial Economics* 33, 3–56.
- Fama, Eugene F., and Kenneth R. French, 2000, "Forecasting profitability and earnings", *Journal of Business* 73, 161–175.
- Graham, John R., Cam Harvey, and Shiva Rajgopal, 2005, "The economic implications of corporate financial reporting", *Journal of Accounting and Economics* 40, 3–73.
- Halpern, Paul, Robert Kieschnick, and Wendy Rotenberg, 1999, "On the heterogeneity of leveraged going private transactions," *Review of Financial Studies* 12, 281–309.
- Henderson, Brian J., Narasimhan Jegadeesh, and Michael S. Weisbach, 2006, World markets for raising new capital, *Journal of Financial Economics* 82, 63–101.
- Jain, Bharat A., and Omesh Kini, 1994, "The post-issue operating performance of IPO firms", *Journal of Finance* 49, 1699–1726.
- Kothari, S.P., Jay Shanken, and Richard G. Sloan, 1995, Another look at the cross-section of expected stock returns, *Journal of Finance* 50, 185–224.
- Lang, Mark H., 1991, "Time-varying stock price response to earnings induced by uncertainty about the time series process of earnings," *Journal of Accounting Research* 29, 229–257.
- Leland, Hayne E., and David H. Pyle, 1977, "Informational asymmetries, financial structure, and financial intermediation", *Journal of Finance* 32, 371–387.
- Lerner, Joshua, 1994, "Venture capitalists and the decision to go public", *Journal of Financial Economics* 35, 293–316.
- Li, Erica X. N., Dmitry Livdan, and Lu Zhang, 2006, "Optimal market timing", Working paper, University of Rochester.
- Loughran, Tim, and Jay R. Ritter, 1995, "The new issues puzzle", *Journal of Finance* 50, 23–51.
- Loughran, Tim, and Jay R. Ritter, 1997, "The operating performance of firms conducting seasoned equity offerings", *Journal of Finance* 52, 1823–1850.

- Loughran, Tim, and Jay R. Ritter, 2004, “Why has IPO underpricing changed over time?”, *Financial Management* 33, 5–37.
- Maksimovic, Vojislav, and Pegaret Pichler, 2001, “Technological innovation and initial public offerings”, *Review of Financial Studies* 14, 459–494.
- Mikkelsen, Wayne H., M. Megan Partch, and Kshitij Shah, 1997, “Ownership and operating performance of companies that go public”, *Journal of Financial Economics* 44, 281–307.
- Pagano, Marco, 1993, “The flotation of companies on the stock market: A coordination failure model,” *European Economic Review* 37, 1101–1125.
- Pagano, Marco, Fabio Panetta, and Luigi Zingales, 1998, “Why do companies go public? An empirical analysis”, *Journal of Finance* 53, 27–64.
- Pástor, Ľuboš, and Pietro Veronesi, 2003, “Stock valuation and learning about profitability”, *Journal of Finance* 58, 1749–1789.
- Pástor, Ľuboš, and Pietro Veronesi, 2005, “Rational IPO waves”, *Journal of Finance* 60, 1713–1757.
- Ritter, Jay R., 1991, “The long-run performance of initial public offerings”, *Journal of Finance* 46, 3–27.
- Ritter, Jay R., and Ivo Welch, 2002, “A review of IPO activity, pricing, and allocations”, *Journal of Finance* 57, 1795–1828.
- Stein, Jeremy C., 1989, “Efficient capital markets, inefficient firms: A model of myopic corporate behavior,” *Quarterly Journal of Economics* 104, 655–669.
- Teoh, Siew Hong, Ivo Welch, and T. J. Hong, 1998, “Earnings management and the long-run market performance of initial public offerings”, *Journal of Finance* 53, 1935–1974.
- Zingales, Luigi, 1995, “Insider ownership and the decision to go public”, *Review of Economic Studies* 62, 425–448.