

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES


CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

IPO POST-ISSUE MARKETS:
QUESTIONABLE PREDILECTIONS BUT DILIGENT LEARNERS?

Peter Bossaerts

Pierre Hillion



SOCIAL SCIENCE WORKING PAPER 1014

August 1997

IPO Post-Issue Markets: Questionable Predilections But Diligent Learners?

Peter Bossaerts

Pierre Hillion

Abstract

Efficiency in the IPO (Initial Public Offering) aftermarket is tested without imposing any restrictions on the priors about potential default at the issue date. Merging Ritter's extended dataset (which covers the period 1975-84) with the CRSP tapes, IPOs are followed up to ten years after issue. Across all IPOs, or when stratifying IPOs according to issue underpricing, industry affiliation or rank of entry in an industry, little evidence against rational price behavior is found. In contrast, the market clearly over-reacts to information about the eventual fate of low-priced issues. A suggestive relationship between irrational price behavior and subsequent takeover activity is uncovered.

Keywords: Initial Public Offering, Market Efficiency, Rational Learning, Default.

JEL Classification: C22, D84, G14.

IPO Post-Issue Markets: Questionable Predilections But Diligent Learners?*

Peter Bossaerts

Pierre Hillion

1 Introduction

Ever since Ritter's seminal empirical study (Ritter [1991]), the post-issue performance of IPOs has been considered to be a puzzle. In the long run, IPOs significantly underperform standard benchmarks or equity in appropriately matched firms. The puzzle has been confirmed in numerous follow-up studies (see, e.g., the Spring 1993 issue of *Financial Management*; Jain and Kini [1994]; Loughran and Ritter [1995]).

The evidence is now generally interpreted as suggesting that the market is too optimistic when pricing young issues. It realizes its mistakes slowly, adjusting prices as the issues mature.

Some have argued that the biases in the market's prior at the issue date are a natural consequence of shortsale restrictions (Miller [1977], Morris [undated]). It could also be a mere sign of the beliefs at a particular point in time. Indeed, most studies focus on IPOs executed the 1970-80s. That priors over this period were biased does not necessarily imply irrationality, because the bias was only demonstrated to be there *ex post*, i.e., with the benefit of hindsight.

Instead, it seems much more fruitful to ask whether subsequent *changes in the market's beliefs* were rational. If beliefs can be expressed in terms of the chance numbers of classical probability theory, we know precisely what this means: changes should obey the rules of conditional probability (Bayes' law).

Examples abound in the experimental literature on *individual* decision making that not everybody uses conditional probability when learning about uncertain events. See, e.g., El-Gamal and Grether [1995]. From the point of view of finance, however, it is

*First author's affiliation: California Institute of Technology; Mailing address: HSS 228-77, California Institute of Technology, Pasadena, CA 91125; Phone (626) 395-4028; e-mail: pbs@rioja.caltech.edu. Second author's affiliation: INSEAD; Mailing address: INSEAD, F-77305 Fontainebleau; Phone +33.1.6072-4214; e-mail: hillion@insead.fr. Many thanks to Jay Ritter for making his dataset available.

important to know whether changes in *market* beliefs, as reflected in prices, obey the rules of rational learning. The goal of the present paper is to test this in the context of IPOs.

A market that fails to update its priors rationally is one that over-reacts or under-reacts to the advent of new information. Changes in prices do not correctly reflect the additional information, in violation of market efficiency (Fama [1970]). Money is to be made: Dutch book arguments could be used to win in betting against a market that violates the rules of conditional probability.¹ In finance terms: *there is an arbitrage opportunity*.

This paper tests the rationality of price changes in the IPO aftermarket with a methodology that requires little or no information about the actual market prior at the issue date and how priors varied across issues. The methodology was originally developed in Bossaerts [1997] and successfully applied to experimental winner-take-all markets in Bondarenko and Bossaerts [1996], to digital option prices implied by index call and put option prices in Bondarenko [1997], and to straight index call options in Bossaerts [1997]. The applications have one thing in common: they concern securities with a clear *bankruptcy state*, like the equity contracts studied here.

The framework of analysis is the following. At the launch date, it is known that a certain number of IPOs eventually fail (default), but it is not known exactly how many and, if the issue at hand does fail, at what time. For simplicity, the recovery rate conditional on bankruptcy is set equal to zero.

Initial priors about the probability of bankruptcy are arbitrary and may vary across IPOs. Price changes in the aftermarket reflect rational updating of these priors from news about the fate of the company. The market is supposed to understand how the news relates to bankruptcy (it knows the conditional probability of the signals given the bankruptcy status). Likewise, the market correctly predicts the future valuation of equity in the company conditional on no default. (If default occurs, the market of course knows that this value will be zero.)

The tests that we use to verify rational updating are simple and powerful. A novelty is that they require the empiricist to split the available sample in a winner category (companies that did not default) and a loser category (companies that defaulted). Standard returns have to be modified slightly and weighted appropriately. If rational updating is rejected, the sign of the statistic provides information on the nature of the inefficiency: whether the market over-reacts or under-reacts to new information.

Our methodology tests for correct updating of priors about the likely default of each company *separately*. The methodology is of the event-study type: each history is mapped in event time, with a common event time zero (the IPO date); one loses potential information from the knowledge that two histories occurred sequentially in calendar time. This

¹Dutch books could be interpreted as repeated forward contracts. See Shervish [1995], p. 654-658, for a technical explanation.

implies, in particular, that our methodology does not investigate whether the default history of companies floated earlier in calendar time were correctly reflected in the priors at the issue date of subsequent IPOs. Our methodology allows there to be such updating (priors can vary arbitrarily across IPO histories), but does not study its rationality.²

Using Ritter's sample (covering U.S. IPOs in the 1975-84 period) and following each IPO up to 10 years after the issue date, we find little evidence of irrationality in aftermarket price dynamics. Looking more closely, however, we find strong evidence of over-reaction in low-priced issues.

The remainder of the paper is organized as follows. The next section discusses the data and summarizes past evidence. Section 3 presents our approach and methodology. Section 4 discusses the aggregate empirical results; Section 5 reports tests conditional on issue information. Section 6 briefly contrasts the results from our methodology with those from a statistic that has traditionally been used to study over-reaction and under-reaction, namely the autocorrelation coefficient. Section 7 concludes. The power of the tests is studied in an Appendix, using a novel model of an inefficient financial market (called the Variable Reversal Delay model).

2 The Evidence

Our results are based on a merging of Ritter's sample of IPOs in the 1975-1984 period and the CRSP 1995 NYSE/AMEX/NASDAQ monthly return tape. The two datasets were merged on the basis of PERM number, or CUSIP when the PERM number was not available.

The following IPOs were eliminated. (i) All the IPOs for which the EXCHANGE variable in Ritter's dataset is equal to 4, corresponding to non-Nasdaq OTC issues. (ii) IPOs for which the CRSP variable SHRCD, the share code, differs from 10 or 11. These are certificates, ADRs, SBI (Shares of beneficial Interest), Units, companies incorporated outside the US, Americus Primes and Scores, closed-end funds, closed-end fund companies incorporated outside the US, and REITS. (iii) IPOs of which the first digit of the SIC code equals 6 or 9, corresponding to financial institutions, insurance, savings and loans (6), and utilities (9).

The latter exclusion was decided on because many IPOs in that category were in fact well-established firms that issued stock on the NYSE/AMEX/NASDAQ because of regulatory changes. The nature of these IPOs differs dramatically from that of the typical one, where a young company is floated, usually in a new area of industrial activity.³

²Loughran and Ritter have argued that priors at the IPO date have recently gotten more in line with the actual post-issue performance, indicating that priors are being updated between two issue dates. See Loughran and Ritter [1995], p. 49.

³We used the SIC code that CRSP assigned to an issue on the 1995 tape to determine exclusion. This code often deviates from Ritter's, which was constructed to better reflect industry affiliation on

The final sample contained 1,856 IPOs, down from 2,609 in Ritter’s dataset. For each IPO, trading and delisting information, as well as a vector of 120 post-issue monthly returns were extracted from the CRSP tape (together with 120 synchronous CRSP equally-weighted and value-weighted index returns). This means that ten-year post-issue performance histories for each IPO were available.

A large fraction of these histories are incomplete, due to delisting. The delisting information is important for the tests to be reported later. CRSP provides delisting codes (DC), which we grouped as follows.

- *Active* issues: DC 100;
- *Mergers*: DCs 200-203;
- *Exchanges*: DC 300-390;
- Liquidations and forced delistings: DCs 550-588, 400 and 700; to be referred to as *Liquidations*;
- Inactive but *Unknown* status: DCs 500-520.

We recorded the delisting status *at the end of each twelve-month period after the issue date*. This status is determined as follows. Consider the end of the j th twelve-month period. If the corresponding point in calendar time is before the recorded delisting date or the issue was never delisted, then DC is set equal to 100 (Active issue). Otherwise, it is set equal to the DC in the CRSP tape.

Table 1 provides descriptive information of the distribution of delisting status for each twelve-month period. Obviously, the number of Active issues decreases as time passes. Only 628 issues have ten-year active trading histories. This does not mean that the other issues went sour: 323 IPOs were merged at one point in the ten-year history and 53 disappeared through an exchange of stock.

Still, 678 IPOs were listed as Liquidations at the ten-year mark. In other words, the probability of failure is rather high (36.5% of the sample). Many of the 174 “Unknown”s should also be considered losers. So, roughly 4 out of 10 IPOs will not be a success.

This rather high failure rate makes all the more dramatic Ritter’s observation that the post-issue performance of IPOs has been dismal. His cumulative average return goes out to three years after the issue date, at which point about one third of the failures (losers) have left the sample, and, hence, do not affect the average returns anymore. Of course, this argument cuts both ways: many mergers and exchanges have also occurred prior to the three-year mark. A large fraction of these generated substantial abnormal

the IPO date itself. Hence, we ended up with some companies that were in the excluded industries at the IPO date. Because they were later re-classified by CRSP, or considered to be in a different industry altogether, these companies were deemed to be more typical IPOs, and, hence, retained in our study.

returns, which will not affect the IPO post-issue performance measure after the merger or exchange.

The IPO average underperformance does not change when extending the horizon beyond Ritter’s three-year window. Figure 1 displays the cumulative average return in excess of the return on the equally weighted CRSP index (CAR-CRSP) up to ten years after the issue date.

We normalized the CAR-CRSP with an estimate of the monthly standard deviation of the mean excess return, multiplied by the square root of 120 (the total length of the return histories, in months). Under the null that the average return in excess of the index equals zero, the path of the normalized CAR-CRSP should be that of a standard Brownian Motion. In particular, its value after t months should be a draw from a normal distribution with mean zero and variance $t/120$.⁴

Figure 1 demonstrates that one can reject the hypothesis that the CAR-CRSP is a standard Brownian Motion. It is most evident from its value at particular points. For instance, at 60 months, the CAR-CRSP is close to -4 , which is far in the left tail of the $N(0, 1/2)$ density, and, hence, generates a very low p -value. Even after ten years, the value of CAR-CRSP is -1.97 , which is significant at the 2.5% level.

The evidence in Figure 1 points to an anomaly: why would common stock IPOs underperform on average a broad, equally-weighted equity index? As mentioned in the Introduction, the puzzle is robust to further scrutiny (alternative performance benchmarks, international evidence, seasoned issues, etc.).

Ritter also studies the aftermarket performance conditional on issue information, such as level of IPO underpricing. This study continues the tradition. In particular, post-issue return histories are stratified according to the following criteria.

- *Underpricing*: issues are sorted by size of underpricing (difference between IPO offer price and first aftermarket price), and then arranged in ten deciles, numbered 0 to 9, with category 0 containing the most underpriced issues.
- *Industry Classification*: issues are arranged in classes depending on the first digit of their three-digit Standard Industry Classification Code at IPO date (as recorded in Ritter’s dataset).
- *Price Level*: issues are sorted by offer price (as recorded in Ritter’s dataset), and then arranged in ten deciles, numbered 0 to 9, with category 9 containing the most expensive issues.
- *Entry Rank*: issues are sorted by order of entry in each industry (determined by the three-digit Standard Industry Classification Code, as recorded in Ritter’s dataset) and arranged in corresponding categories. Category 1 includes the IPOs that were

⁴This follows from a standard functional central limit theorem.

recorded to be the first (in the dataset) in their respective industries. Category 2 contains all second-comers. Etc.

With the Underpricing categorization, one can investigate whether the IPO anomaly is specific to extremely underpriced issues. Likewise, the Industry Classification, Price Level and Entry Rank categorizations enable one to relate aftermarket underperformance to industry and offer price, respectively, or to study whether the anomaly is specific to early issues in an industry.

3 Theory

3.1 Mere Optimism Or Fads?

Ritter's findings have been interpreted as evidence of optimism, or even fads. The term *fads* reflects something irrational: the market overprices the issues, and stubbornly holds on to its beliefs, until finally correcting in the face of evidence of failures.

In fact, the optimism is not necessarily irrational. Optimism is merely a characterization of someone's beliefs at a point in time. Optimism becomes irrational only if it is not corrected (updated) properly as contrary evidence emerges. So, optimism is not the problem; instead, the obstinacy is.

Rationality is a property of learning, and not of beliefs. The aim of this paper is to investigate whether IPO underperformance is merely a consequence of optimistic beliefs at the issue date of the IPOs in Ritter's dataset, or whether it also reflects obstinacy (or its opposite, overexuberance). The latter would be evidence of irrationality.

The paper does so in the context of a framework which allows the empiricist to be agnostic of the actual priors of the market at the issue date and to avoid specification of the payoff generating process beyond the proposition that there is a bankruptcy state in which the issue becomes worthless.

The methodology was originally developed in Bossaerts [1997] and will be summarized here in terms of the IPO application. The framework is presented first; the tests are discussed next.

3.2 Framework

Let t index (event) time, measured in, say, months. We want to test rational learning in a period up to T months after the issue date, based on the status of the IPO at some future reference point $T^* > T$. T is a fixed point in event time, e.g., 18 months after the issue date, or T is the date of the last available price if delisting occurred prior to this fixed point.

Assuming that no dividends are paid, or that dividends are continuously reinvested, the price at t , p_t , predicts the value of equity (with dividends reinvested) in the future. Let V_{T^*} denote the value of equity at the future reference point T^* . Let $R_{F,t}$ be the (deterministic) riskfree rate in month t . Let ϕ_t denote the market's information at time t . Assuming risk neutrality,

$$p_t = \frac{1}{(1 + R_{F,t+1}) \dots (1 + R_{F,T^*})} E[V_{T^*} | \phi_t]. \quad (1)$$

The conditional expectation in this expression is computed on the basis of the market's subjective beliefs, which need not coincide with the actual distribution of V_{T^*} . It will be specified shortly to what extent the market's beliefs may be incorrect.

Risk neutrality and deterministic interest rates will be assumed throughout. It should be emphasized that the tests reported below could be altered to accommodate stochastic interest rates, or even risk aversion. See Bossaerts [1997]. Nevertheless, we will be able to report that our theory provides a clean and parsimonious interpretation of post-IPO performance while sticking to risk neutrality and deterministic interest rates.

By time T^* , the issue may have defaulted, in which case the recovery rate is assumed to be zero, and, hence, $V_{T^*} = 0$.⁵ Absent bankruptcy, $V_{T^*} > 0$. The market may hold incorrect beliefs about the frequency of occurrence of bankruptcy. Conditional on no default, however, the market's beliefs about the value V_{T^*} are correct.

The actual bias in the market's prior about bankruptcy are left unspecified. It can be arbitrary. Moreover, it can vary across issues. We do require, however, that the market learn rationally about the potential bankruptcy of individual issues, using its information ϕ_t . In other words, the market uses the rules of conditional probability (Bayes' law) to update its beliefs.

It was mentioned that the position of the reference date T^* is arbitrary: any date after T would do. It must lie *after* T , however ($T^* > T$). This is because we cannot allow for perfect revelation that the company defaulted during the period for which returns are computed ($t = 1, \dots, T$). Part of our analysis will require us to compute returns with the *end-of-period price as basis*. These risk to be zero if default is revealed to have occurred, making the corresponding return measure unbounded. This causes obvious statistical problems.⁶

⁵It is possible to accommodate a nonzero recovery value. It is necessary, however, that the recovery value be *fixed*. The test results that are reported in this paper should be adjusted slightly if the recovery value is argued to be nonzero.

⁶See Bossaerts [1997] for an in-depth analysis. In the IPO study, many companies were known to have defaulted prior to our reference date T^* . This does not invalidate our tests, however: the important restriction is that knowledge of default is not reflected in the price history we use to compute returns, i.e., prices at dates $t = 1, 2, \dots, T$. Delisting announcements of distressed firms invariably occur *after* the last available price on the CRSP tape, T . Hence, T is always before potential bankruptcy revelation. For a related point, see Shumway [1997].

In our world, it will no longer be true that changes in prices (properly deflated) will be unpredictable. Mathematically,

$$E^* \left[\frac{1}{1 + R_{F,t+1}} p_{t+1} - p_t \mid \phi_t \right]$$

may be nonzero. In contrast with (1), the conditional expectation is now computed from the actual (true) distribution of price changes over month $t + 1$ across issues (an asterisk makes this distinction clear). Defining the excess return to be

$$r_{t+1} = \frac{\frac{1}{1 + R_{F,t+1}} p_{t+1} - p_t}{p_t},$$

this can be translated to mean that the (excess) return may be nonzero on average. Mathematically, the following *need not hold*:

$$E^* [r_{t+1} \mid \phi_t] = 0. \tag{2}$$

Yet, extant studies of IPOs have focused on testing whether (2) holds. The tests reject. Rejections occur even when adjusting returns for risk by, e.g., subtracting the (excess) return on a broad index.

3.3 Methodology

We allow beliefs about the probability of bankruptcy to be arbitrary, and even to vary across issues. Moreover, we remain agnostic about the processes generating payoffs and information. In such a context, one would reasonably conjecture that rational learning does not impose falsifiable restrictions. In other words, any set of price histories could be explainable in terms of some set of biases in the market's prior and some payoff and information generating processes.

Bossaerts [1997] proves that this conjecture is wrong. Rational learning does restrict the dynamic behavior of prices, even in this fairly generic environment. The restrictions are not apparent from a study of excess returns as in extant studies of IPOs.

Foremost, the restrictions emerge only after splitting the sample in two subsamples, one of histories of *winners* (issues that did not default by T^*) and those of *losers* (issues that defaulted by T^*). In fact, this is a virtue: it directly addresses an issue that plagues the interpretation of extant empirical results. The abnormal behavior of the (normalized) cumulative average return in Figure 1, for instance, could be attributed to complex selection biases, because the composition of the cross-section on which the performance measure is based changes with event time.

Further, the selection bias introduced by studying the subsamples of winners and losers separately implies that (excess) returns will not necessarily be equal to zero, even

if priors happen to be correct. Indeed, excess returns on winners, for instance, can be expected to be positive on average. As a consequence, performance measures that differ from the traditional return will have to be studied.

Here are the main restrictions.

First, consider the winners, i.e., IPOs that did not default before T^* . Define the *modified* (excess) return:

$$\begin{aligned} x_{t+1} &= \frac{\frac{1}{1+R_{F,t+1}}p_{t+1} - p_t}{\frac{1}{1+R_{F,t+1}}p_{t+1}} \\ &= \frac{p_{t+1} - (1 + R_{F,t+1})p_t}{p_{t+1}}. \end{aligned} \tag{3}$$

The modified (excess) return differs from the traditional return in that the *future* (deflated) price is used as basis.

For W issues that did not default before T^* (indexed $i = 1, \dots, W$), compute the *average modified excess return*:

$$\frac{1}{W} \sum_{i=1}^W \left(\frac{1}{T-1} \sum_{t=1}^{T-1} x_{i,t+1} \right). \tag{4}$$

Fact 1 *For winners, the expected average modified excess return is nonpositive. If the winners' payoff is random, the expected average modified excess return is strictly negative.*

To build some intuition as to why this Fact is true, consider the average traditional (excess) return. If this is computed on the basis of histories of winners only, one expects a positive bias, at least (and this is important), *if the market correctly reacted to news that the issues were going to be winners*. To offset this bias, positive (excess) returns should be multiplied by a factor smaller than one, and negative (excess) returns should be multiplied by a factor larger than one. The ratio $p_t/(p_{t+1}/(1 + R_{F,t+1}))$ is such a factor. Multiplying the traditional (excess) return with this factor produces the modified (excess) return. Figure 2 provides a graphical illustration. It turns out that the chosen factor more than offsets the bias; the resulting variable, namely, the modified (excess) return, will be nonpositive on average. It will be *strictly negative* if the payoff is random.

Fact 1 is an implication of rational learning. In particular, the strict negativity of the modified (excess) return obtains not just because the market reacted favorably to news that the issues at hand eventually became winners, but precisely because this reaction was *correct*, i.e., in accordance with Bayes' law.

It deserves emphasis that market priors about the likely default of the issue may be arbitrary; in particular, they need not be correct, and they can vary across issues.

A specific weighting scheme makes the average modified excess return *zero*. Let V_{i,T^*} denote the time- T^* value on winner issue i ($i = 1, \dots, W$). Consider the following *weighted average modified excess return*:

$$\frac{1}{W} \sum_{i=1}^W V_{i,T^*} \left(\frac{1}{T-1} \sum_{t=1}^{T-1} x_{i,t+1} \right). \quad (5)$$

Fact 2 *For winners, the expected weighted average modified excess return is zero.*

Again, the intuition is simple. The weighting scheme gives more weight to issues with a large payoff at T^* . These are the issues that must also have had the larger price run-up, i.e., higher (modified excess) returns, in anticipation of the larger payoff. Indeed, we assume that market prices correctly incorporate the generally favorable information. Issues with lower time- T^* values have had lower (modified excess) returns, but are weighted less in the weighted average modified excess return. Altogether, the weighted average modified excess return is expected to be zero.

Again, it should be underscored that Fact 2 obtains irrespective of the market's prior about default: it may be biased and it may vary across issues.

Some may find Fact 1 to be redundant, in view of the cleaner restriction in Fact 2. Fact 2, however, uses a weighting which may affect the power of the test, depending on the data at hand. We will discuss the issue of power shortly.

The restrictions in Facts 1 and 2 apply only to modified returns computed for winners. Consider next the losers, i.e., issues that defaulted before T^* .

Here we need a restriction on the distribution of the market's prior about likely default across issues. In particular, it must be diffuse (non-informative). In other words, the market's prior about default must not be fixed and its distribution must not be biased in a particular direction (optimism, pessimism). The fact that the prior must not be fixed implies that it cannot be correct either.

For L losers, indexed $i = 1, \dots, L$, compute the *average excess return*:

$$\frac{1}{L} \sum_{i=1}^L \left(\frac{1}{T-1} \sum_{t=1}^{T-1} r_{i,t+1} \right). \quad (6)$$

Fact 3 *For losers, the expected average excess return is zero if the market's prior about default is diffuse across IPOs.*

This Fact follows from a study of the behavior of securities prices in *reverse time*. In particular, consider the question whether one can back-predict prices of losing issues. The answer turns out to be negative when the distribution of the market's prior is diffuse. For if this distribution were informative because, say, priors are fixed at the correct level, then a low price level would indicate that prices must have decreased from an average

(correct) initial level. Hence, returns could be back-predicted to have been negative on average.

If returns cannot be back-predicted from the price level, their *unconditional* average must also be zero, by the law of the iterated expectation. This is what Fact 3 states.

Fact 3 leads to a *joint test* of the proposition that the market updates its priors about default rationally *and* that the distribution of the priors is diffuse across IPOs. Rejections may merely reflect violation of the latter, and not of the former.

The three Facts were expressed in terms of unconditional averages. Facts 1 and 2 also obtain if *conditioning on issue information*, such as IPO underpricing or IPO offer price.

Fact 4 *Facts 1 and 2 continue to hold when the averages are computed for subclasses of winning IPOs formed on the basis of issue information.*

If the averages are not computed from the first aftermarket trade on, but from some future initial date $\tau < T$, Facts 1 and 2 also hold when conditioning on information that emerged after the issue date, up to time τ .

The four Facts can easily be tested, because they restrict the behavior of readily available statistics, based on simple price information. We chose not to use absolute price information, but *wealth ratios*, based on an investment at IPO of \$1, with reinvestment of possible dividends. That way, the future values V_{i,T^*} become normalized, reducing their cross-sectional variation, and improving the power of tests of Fact 2.

We also decided to split the winner and loser categories into subclasses. Winners were stratified into Active, Mergers and Exchanges, in accordance with the classification explained in the previous section. Likewise, losers were subdivided into Liquidations and Unknown.

The Facts do not necessarily hold for these subclasses, which are formed on the basis of future information *beyond* the default status. Nevertheless, the rejections of some of the Facts for some subclasses would be informative about certain issues of corporate finance. For instance, if we only discover violations of Fact 2 in the Mergers category, we could be lead to believe that mergers occur *because* of irrational price behavior. We will elaborate on this point when we discuss the results.

The reader may have wondered about the asymmetry in the results for winners and losers. This is due to the asymmetry of the framework: winners pay a random amount, losers pay a fixed, known quantity (zero). Symmetry re-emerges when winners pay a fixed amount as well, as in the case of digital options, which pay \$1 if they mature in the money. There, the weighted average modified return coincides with the (unweighted) average modified return. Hence, Facts 1 and 2 together imply that the average modified return of winning digital options will be zero if learning is rational. For digital options that mature out of the money, the traditional return must be zero (Fact 3). But the

traditional return becomes a modified return when expressed in reverse time. Since the result for losers is obtained from a reverse-time analysis, it should actually be stated as implying that the modified return in reverse time should be zero. Altogether, the average modified return forward in time will be zero for winners; the average modified return in reverse time will be zero for losers. Symmetry is restored.

3.4 Meaning Of Rejections

Before turning to the results, it is important to delineate the information that would be conveyed by *rejections* of the Facts.

Rejections of Fact 3 could merely reflect violations of the assumption that the distribution of the market's prior about bankruptcy is diffuse. When the market's prior is *correct*, as has been required in tests of the rationality of post-IPO price performance in extant research, one expects there to be violations. In fact, unless one is convinced that the distribution of the market's prior is diffuse, one would hope to discover rejections; this would indicate that the test has power.

Most information could be conveyed by tests of Fact 2, because, unlike Fact 1, the alternative hypothesis would include two cases: positive and negative weighted average modified excess returns. Here is an interpretation in terms of a market whose price reaction may be correct in the long run, but either too slow or too excessive in the short run.

Remember the intuition behind Facts 1 and 2: the traditional (excess) return on winners is positive, because of the selection bias and the market's learning; to offset this, (excess) returns must be multiplied by a factor which is below one for positive returns and above one for negative returns; if the market learns rationally, the factor that does the job is the ratio of today's price over tomorrow's (deflated) price; it converts the standard (excess) return into the modified (excess) return.

Positive weighted average modified excess returns are a sign of *under-reaction* of the market to new information. In other words, price changes amount only to a fraction of the full Bayesian update; further adjustments are made later. It implies that price movements are less extreme than in the rational case, making the factor that transforms standard returns into modified returns less variable. As a result, the bias caused by looking only at winners is not entirely offset.

Negative weighted average modified excess returns are a sign of *over-reaction* of the market to new information. Price changes overcompensate for the Bayesian update, creating the necessity for future reversals. It implies that price movements are more extreme than in the rational case, making the factor that transforms traditional returns into modified returns more variable. Hence, the selection bias is over-adjusted.

Figure 3 displays three price paths for a digital option that matures in the money.

A digital option either pays \$1 or \$0. Because the payoff in the no-default state is fixed and known (\$1), the weighted average modified return and the average modified return coincide, and must both be zero if prices reflect correct updating. The solid line represents such a case: the average modified return equals 0.05%. In an over-reacting market, prices are more volatile than warranted by rational price adjustment. The dashed line depicts such a case. The average modified return is negative and almost two orders of magnitude larger (-1.06%). The dotted line represents under-reaction: prices adjust too slowly to new information. The average modified return is positive (1.02%).

In classical tests of market efficiency, return autocorrelations convey signals of over-reaction and under-reaction of the market to new information. If priors may be biased, return autocorrelations do not provide the right signal. Moreover, autocorrelations only provide information about the behavior of *deviations* of returns from the mean; they do not test the appropriateness of the *mean* return.

In contrast, the weighted average modified return of winners must be zero even if priors are biased (Fact 2). Moreover, it conveys information about the correctness of both price changes net of the average change, *and* the average change itself, as the above discussion points out.

Because of their popularity in tests of market efficiency, we will come back to return autocorrelations at the end of the presentation of the results. We will document return autocorrelations for IPOs. Their message is confusing, in contrast to the conclusions we draw from tests of our Facts.

In the Appendix, we construct a simple model of an over-reacting or under-reacting market. The construction is the basis for a Monte Carlo analysis of the power of tests based on Fact 2. It also reveals that autocorrelations may be low even in a substantially irrational market. In contrast, tests based on Fact 2 have power.

4 Unconditional Tests

We first report results from unconditional tests (Facts 1 to 3), for winners and losers separately.

4.1 Winners

Figures 4 to 6 provide graphical evidence about Facts 1 to 3 for each of the winner subcategories. The plots provide point estimates and 95% confidence intervals.

The Figures report averages for different post-issue reference points (T^* in the previous section), spaced at twelve-month intervals. For each reference date, averages were computed using return data from the first after-market trade up to six months before

the reference date at the latest. (In terms of the notation of the previous Section, T is positioned up to six months before T^* .)

IPOs were followed up to ten years after issue date; hence, there were ten reference points, at the end of months 12, 24, ..., 120. Reference point values (V_{i,T^*} in the previous section) for issues that merged or were exchanged (Mergers and Exchanges) were computed by riskfree reinvesting all the wealth accumulated up to the date of the last available return on the CRSP tape.

The three-month Treasury bill rate was used throughout for riskfree discounting.

The averages displayed in the Figures are computed on the basis of overlapping periods (the averages for the period ending after two years are based on observations from the first year as well). Hence, the results are not independent across reference points. This continues a tradition of the event-study literature in general and the empirical IPO literature in particular.

In addition to the average modified excess return, the average (traditional) excess return and the weighted average modified excess return, the Figures display results for the weighted average *traditional* excess return. This statistic is computed from (5) after substituting the traditional excess return $r_{i,t+1}$ for the modified excess return $x_{i,t+1}$. It gives a good benchmark to gauge the power of the test. (Of course, the presence of rejections in itself will confirm power.) We expect the average traditional returns, weighted or not, to be strictly positive, because of the bias from winner selection. The power of the tests should be questioned, however, if the weighted average excess return is seldom found to be significantly positive.

Our reporting the (equally weighted) average traditional return in the Figures plays the same role: one expects it to be significantly positive; if not, there may be too much noise in the data for our tests to be powerful.

Fact 1 is *overwhelmingly accepted*. In the case of Active issues (by far the most numerous subclass), we reject for all future reference points that the average modified excess return is non-negative. In contrast, the average traditional excess return is significantly positive for all reference points beyond two years.

There are few mergers in the first year after an issue (see Table 1), which explains why the average modified excess return is significantly negative only from the second year on. Likewise, the positive selection bias shows up in the average traditional excess return only beyond the first 5 years.

There are even less observations on exchanges, which explains why the average modified excess return is significantly negative only from the sixth year on, and the average traditional excess return becomes significantly positive only after five years.

The restriction of Fact 2, that the weighted average modified excess return ought to be zero, is *never rejected*. Contrast this with the weighted average traditional excess return,

which is *always* significantly positive (despite the variability of the weights), except when few observations are available (first year for Mergers; first and second year for Exchanges).

Based on the evidence from the unconditional tests, it is fair to conclude that the post-issue price series do not reflect any irrationality. Priors may have been biased, but this did not keep the market from correctly learning about future values.

4.2 Losers

Figure 7 displays evidence against Fact 3 for the two loser categories, Liquidations and Unknown: the average (traditional) excess return is almost invariably significantly negative.

It is not clear whether this reflects irrational price adjustment in the loser family. As mentioned in the previous section, Fact 3 requires the distribution of market priors to be diffuse across issues. If priors are fixed (which certainly would be the case if the market held correct beliefs), one expects to see violations of Fact 3.

5 Conditioning On Issue Information

We now report results on price histories sampled on the basis of information from the IPO itself. This should verify Fact 4. We will focus on the weighted average modified excess return, which is the most convenient and cleanest statistic to gauge the rationality of price behavior in the aftermarket.

5.1 Underpricing

Figure 8 demonstrates that there is no evidence of irrationality in any specific subcategory of Active issues based on issue underpricing, confirming the unconditional evidence. With few exceptions, the weighted average modified return is insignificantly different from zero, confirming Fact 2.

There is some evidence against Fact 2 in decile 7 of the Mergers and deciles 8 and 9 of the Exchanges (although the latter evidence must be treated with caution because of scarcity of observations). See Figures 9 and 10. The lack of any specific relationship between underpricing decile and evidence against Fact 2 (why decile 7 of the Mergers?) leads us to conclude that the evidence against rational aftermarket price behavior is dubious.

5.2 Industry Code

Using Ritter’s industry classification, we find some systematic, industry-specific evidence against Fact 2. See Figures 11 to 13. Industry 1 (with Standard Industry Classification Code 1**) generates significantly negative weighted average modified excess returns across the three winner categories. This appears to indicate that the IPO aftermarket in industry 1 *over-reacts* to new information.

For Exchanges, evidence of over-reaction also emerges in industries 4 and 5, and of *under-reaction* (significantly positive weighted average modified excess returns) in industry 8. Nevertheless, the scarcity of observations cautions us to draw strong conclusions.

5.3 Initial Price Level

The most solid evidence against Fact 2 emerges from low-priced issues. Figures 14 to 16 displays ample rejections for the three lowest deciles, and some scattered rejections in deciles 4 to 6. In all the rejections, the weighted average modified excess return is negative, which would imply that the aftermarket *over-reacts to new information*.

It is remarkable that issue price level would have anything to do with irrationality in the aftermarket. One usually suspects issue underpricing to be an indicator of irrationality (“hot issues”). The evidence from the weighted average modified excess return refutes this. In contrast, the offer price predicts aftermarket anomalies.

This is not unlike the finding of over-reaction in changes in the market value of low-priced call options written on the S&P 500 index. See Bossaerts [1997]. Low-priced call options are almost invariably out-of-the-money. They are the equivalent of a highly levered position in the underlying security, i.e., an inexpensive bet. Future research should indicate whether low-priced IPOs are similarly perceived to be cheap bets.

The solid rejections are good for the methodology that is used in this paper: they demonstrate that the tests have power. They are bad for the status of market efficiency in the IPO market, however, because they imply that the aftermarket may not only hold biased priors: they also reveal that the market does not even know how to rationally react to new information.

5.4 Entry Rank

Figures 17 and 18 depict results for Active issues and Mergers as a function of their rank of entry in an industry. Only results for the first ten entrants are reported. Also, results for Exchanges are not reported because of lack of observations.

Because the beginning point of the Ritter dataset is rather arbitrary, a categorization on the basis of entry rank may not be meaningful, and, hence, produce uninformative

results. The fact that disproportionately few early entrants lead to exchanges, however, indicates that the classification may not be arbitrary after all.

There is little evidence against Fact 2 from early entrants in the Active category. Some evidence in favor of over-reaction (significantly negative weighted average modified excess returns) emerge for later entrants; while not reported, this pattern continues all the way to the 20th entrant. The finding would not be inconsistent with a world in which new industries generate substantial interest of “irrational traders” once the first few IPOs are consummated.

Another interesting observation emerges from the significantly *positive* weighted modified excess returns for the first entrants that become merged at the end of the respective one-year periods. These are not an artefact of a few observations (by the tenth year, 88 first-entry IPOs fall into the Merger category).

The positive weighted average modified excess returns would reflect *under-reaction*: the market does not fully update on the basis of the available information. The fact that evidence of under-reaction surfaces in the Merger category is intriguing: the eventual merger itself may have been the consequence. This brings us to a more general point about takeovers.

5.5 The Role Of Takeovers

Much of the evidence against Fact 2 and its premise of rational aftermarket price behavior comes from issues that eventually merged or were exchanged. The takeover may not have been incidental to the irrational price behavior.

The nicest example comes from the positive weighted average modified excess returns for first entrants in an industry, discussed in the previous Subsection. The evidence suggests that under-reaction may have caused a takeover. If under-reaction is tantamount to undervaluation, the implications for the beneficial role of mergers, targeted at irrational price behavior, are obvious. This observation certainly warrants future research.

6 Information In Autocorrelations

It was mentioned before that return autocorrelations may give a confusing or even misleading signal about the rationality of a market. We now illustrate this.

Figure 19 plots the autocorrelations up to lag six (averages across issues) for *all* the IPOs (no selection bias induced). The issues are divided into deciles depending on offer price. The autocorrelations are generally negative but small. They are often significant, however. No particular pattern can be discerned: high-priced issues seem to generate equally significant autocorrelation coefficients.

It is not clear whether these autocorrelations reflect rational adjustment of biased initial market priors, or outright irrational price behavior (in particular, over-reaction). Moreover, the pattern is confusing, and many would consider the level of autocorrelation to be economically insignificant (which may be the reason why previous empirical studies of IPOs have not dwelled upon this issue).

Our study can be understood as providing an explanation for the autocorrelation patterns (as well as the low average performance of IPOs overall). It exemplifies that autocorrelations are not necessarily informative.

We now re-introduce a selection bias and report autocorrelations for Active issues and Mergers separately, in Figures 20 and 21, respectively. The autocorrelations only partially pick up the irrationality in the aftermarket behavior of low-priced issues. More confusingly, they would also suggest anomalies in high-priced issues. The weighted average modified excess returns displayed in Figures 14 and 15 prove that this reading is wrong.

7 Conclusion

In the aggregate, we find little evidence against rational aftermarket price behavior for U.S. IPOs between 1975 and 1984. If, with hindsight, initial aftermarket price levels are found to be incorrect, it must be the consequence of optimism at the time of the IPO. There is certainly no evidence that the market over-reacted or under-reacted as news about the fate of each individual issue emerged.

At a more microscopic level, however, we find solid evidence against rational price behavior in the aftermarket of low-priced issues. The results mostly point to over-reaction. We also discover a relationship between aftermarket irrationality and subsequent merger or stock exchange activity.

Throughout, we have not adjusted for risk. The assumption of risk neutrality has been maintained. We were able to organize and explain the data without having to appeal to models of risk premia. This is attractive, for two reasons. Models of risk premia are controversial; also, for reasonable levels of risk aversion, risk premia are of second order, and, hence, one can question their potential to explain actual excess returns.

In our replication of Ritter's results, however, we computed IPO returns in excess of the return on the equally weighted CRSP index. When we plot returns in excess of the riskfree rate (which is the appropriate benchmark under risk neutrality), we find a significant *overperformance* at the ten-year reference point. See Figure 1. How, then, should we interpret this finding?

As mentioned when discussing Figure 1 earlier on, one ought to be cautious because of selection biases: as time progresses, the performance index is subject to a severe

survivorship bias. In particular, a large fraction of distressed firms quit the sample. In itself, this could explain the outperformance.

Absent selection biases, we would still not insist on zero outperformance, however, because an integral part of our approach has been to allow the market to have biased priors at the issue date. In other words, the theory of market efficiency on which our tests are based does not imply that average returns in excess of the riskfree rate must be zero. Plots of (normalized) cumulative average returns could go in any direction. Figure 1 confirms this.

If one is unwilling to ignore the selection biases that plague the interpretation of plots like Figure 1: our methodology addresses this problem directly. Indeed, the tests *require* the very selection biases that cause the unease. In particular, returns on winners and losers have to be analyzed separately.

Let us state this differently. The interpretation of standard tests of efficiency of the IPO aftermarket is rendered difficult, if not impossible, by complex selection biases. Which implies, in particular, that one should not jump too quickly to conclusions about the biases in the market's priors at the issue date. In contrast, this paper demonstrates that market efficiency can be tested without requiring that priors be unbiased. Stronger even, the tests allow one to remain agnostic about the nature of priors. They better do: efficiency is a proposition about whether the market incorporates information correctly in prices, i.e., whether it learns rationally, and not about priors. Only methodological constraints hitherto forced the empirical finance literature to confound the two.

The methodology in the present paper could be applied to event studies of other financial markets, such as the markets for corporate or sovereign bonds. In general, it can be used to study the prices of any financial contract that features a clear "default state," in which a fixed and known payout occurs. This includes straight call options, as illustrated in Bossaerts [1997].

References

- Bondarenko, O., 1997, "Testing Rationality of Financial Markets: An Application to S&P 500 Index Options," Caltech working paper.
- Bondarenko, O. and P. Bossaerts, 1996, "Expectations and Learning in Iowa," Caltech working paper.
- Bossaerts, P., 1997, "The Dynamics Of Equity Prices In Fallible Markets," Caltech working paper.
- El-Gamal, M. and D. Grether, 1995, "Are People Bayesian: Uncovering Behavioral Strategies," *Journal of the American Statistical Association* 90, 1137-45.
- Fama, E., 1970, "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance* 25, 383-417.
- Jain, B. and O. Kini, 1994, "The Post-Issue Operating Performance of IPO Firms," *Journal of Finance* 49, 1699-1726.
- Loughran, T. and J. Ritter, 1995, "The New Issues Puzzle," *Journal of Finance* 50, 23-51.
- Miller, E., 1977, "Risk, Uncertainty and Divergence of Opinion," *Journal of Finance* 32, 1151-68.
- Morris, S., undated, "The Re-Sale Premium For Assets In General Equilibrium," Department of Economics, University of Pennsylvania, working paper.
- Ritter, J., 1991, "The Long-Run Performance of Initial Public Offerings," *Journal of Finance* 46, 3-27.
- Schervish, M., 1995, *Theory of Statistics*. New York: Springer Verlag.
- Shumway, T., 1997, "The Delisting Bias in CRSP Data," *Journal of Finance* 52, 327-40.

Appendix

To assess the power of tests based on Fact 2, we simulated the following pricing model.

Consider three processes W_t , W_t^n , W_t^d , $t = 0, 1, 2, \dots, T, T + 1, \dots, T^*$. We set: $W_0 = 1$ and $W_0^n = W_0^d = 0$. Consider a security that pays $W_{T^*} - 1$ at T^* if $W_{T^*} > 1$ and zero otherwise. This is a standard call option with strike price 1. “Default” corresponds to the state when $W_{T^*} \leq 1$. In the notation of the paper,

$$V_{T^*} = \begin{cases} W_{T^*} - 1 & \text{if } W_{T^*} > 1; \\ 0 & \text{if } W_{T^*} \leq 1. \end{cases}$$

Under risk neutrality and zero discounting, the time- t “rational” price of the call equals:

$$p_t = E[1_{\{W_{T^*} > 1\}}(W_{T^*} - 1) | W_t]$$

($1_{\{\cdot\}}$ denotes the indicator function.)

To study the power of our test of Fact 2, we propose the following model of an inefficient (over-reacting) market. At each point in time, the actual price, \tilde{p}_t , equals p_t , plus a noise term ν_t , which reflects the market’s irrationality. In particular,

$$\tilde{p}_t = \max(\epsilon, p_t + \nu_t),$$

where ϵ is some small, positive number, and

$$\nu_t = \delta_n \sum_{\tau=1}^t 1_{\{t-\tau \leq \min(\text{round}(|W_\tau^d|), 10)\}} (W_\tau^n - W_{\tau-1}^n);$$

“round” denotes the rounding function, and δ_n is a scaling factor. (The lower bound on prices was chosen because call option prices have to be non-negative, otherwise there is an arbitrage opportunity; the lower bound, ϵ , must be strictly positive, in order to avoid unbounded modified returns when (noisy) prices drop to zero.)

The noise term adds temporary shocks to the rational price. The time that the shock $W_\tau^n - W_{\tau-1}^n$ at τ influences the market price is determined by W_τ^d , but will not exceed ten periods.

We will refer to this noise model as the *Variable Reversal Delay* (VRD) model.

Assume that the increments of W_t are those of a geometric Brownian Motion with drift 0 and variance $(0.1)^2$. If the market knows this (which implies that it has unbiased beliefs about the probability of default), then:

$$p_t = W_t N(d_1) - N(d_2),$$

where $N(\cdot)$ denotes the normal distribution function, and

$$d_1 = \frac{\ln W_t}{0.1\sqrt{T^* - t}} + (0.5)(0.1)\sqrt{T^* - t}; \quad d_2 = d_1 - 0.1\sqrt{T^* - t}.$$

Furthermore, we specify the increments of W_t^n to be independent $N(0, (0.05)^2)$ draws, and those of W_t^d to be $N(0, 1)$ draws. To accommodate varying levels of irrationality, we set the scaling factor δ_n equal to 0 (no irrationality), 0.1, 0.2, ..., 0.5. Because the standard deviation of increments of W_t^n is half that of the signal W_t , most of the variability in prices will still be related to rational updating, even with a scaling factor of 0.5.

When the scaling factor δ_n is as low as 0.1, the VRD model generates only small deviations from rational pricing. This is demonstrated in Figure 22, which displays one price path ($t = 1, \dots, T$, with $T = 80$) for a call option that matured in the money (it paid \$0.35 at $T^* = 100$). The deviations may seem inconspicuous, but they are highly persistent. One wonders whether our test would extract power out of the persistence. We will be able to validate this conjecture.

We deliberately use small sample sizes in the simulations. This will enable us to gauge the small-sample properties of our test if prices are not disturbed by the VRD model. Below, we report results from sets of 50 simulations of 50 sample paths (expect only half of them to lead to winners!). $T = 80$ and $T^* = 100$.

For each simulation of 50 sample paths, we computed the weighted average modified return. When δ_n equals 0 (pricing is rational), Fact 2 states that this statistic is zero on average. For nonzero δ_n s, we expect *negative* weighted average modified returns, because our VRD noise model adds over-reaction to the rational price. See the discussion in Section 3.4.

Figure 23 displays notched boxplots of the distribution of the weighted average modified return for the six sets of simulations. As the noise level increases, the distributions are shifted downward. With zero noise, the mean of the weighted average modified return is -0.1313. With a standard error of 0.1721, this is insignificantly different from zero, confirming Fact 2.

With a level of noise as low as 0.1, the mean of the weighted average modified return is already highly significant (with a p value less than 0.001): it equals -0.6778, and its standard error is 0.1328. This result confirms that small deviations from rationality can lead to rejections, as long as they are persistent, as in Figure 22.

When pricing is rational ($\delta_n = 0$), one observes negative skewness in the distributions of the weighted average modified return, because of the small sample sizes and because modified returns are bounded from above by +1, but unbounded otherwise. This skewness affects inference: there is a higher probability of rejecting the null in favor of over-reaction than suggested by a test based on the normal distribution.

Table 1
Distribution of IPO status at various points up to ten years after issue date

Reference Point	Active	Mergers	Exchanges	Liquidations	Unknown
12 Months	1809 (97.5)	10 (0.5)	1 (0.1)	21 (1.1)	15 (0.8)
24 Months	1627 (87.7)	48 (2.6)	7 (0.4)	109 (5.9)	65 (3.5)
36 Months	1442 (77.7)	81 (4.4)	13 (0.7)	199 (10.7)	121 (6.5)
48 Months	1274 (68.6)	122 (6.6)	24 (1.3)	288 (15.5)	148 (8.0)
60 Months	1113 (60.0)	167 (9.0)	36 (1.9)	378 (20.4)	162 (8.7)
72 Months	965 (52.0)	212 (11.4)	45 (2.4)	464 (25.0)	170 (9.2)
84 Months	848 (45.7)	247 (13.3)	50 (2.7)	541 (29.1)	170 (9.2)
96 Months	764 (41.2)	275 (14.8)	50 (2.7)	595 (32.1)	172 (9.3)
108 Months	689 (37.1)	295 (15.9)	51 (2.7)	649 (35.0)	172 (9.3)
120 Months	628 (33.8)	323 (17.4)	53 (2.9)	678 (36.5)	174 (9.4)

Remarks: The results are based on a merger of Ritter's 1975-84 sample and the CRSP 1995 NYSE/AMEX/NASDAQ monthly return tape; Categories: (i) *Active* issues: CRSP delisting code (DC) 100; (ii) *Mergers*: DCs 200-203; (iii) *Exchanges*: DCs 300-390; (iv) *Liquidations* (and forced delistings): DCs 550-588, 400 and 700; (v) *Unknown* (and inactive): DCs 500-520; Total sample size: 1856; Numbers in brackets: percentage of total.




Figure 1: Normalized cumulative excess returns up to ten years after issue date, U.S. IPOs in the period 1975-84. Under the null hypothesis of an efficient market with unbiased beliefs, and barring selection biases, normalized cumulative excess returns should form the path of a standard Brownian Motion. In particular, their distribution at month t is $N(0, t/120)$.




Figure 2: Top panel: when prices increase, the modified return (based on the end-of-period price level) is lower than the traditional return (based on the beginning-of-period price level). Bottom panel: when prices decrease, the modified return is larger (in absolute value) than the traditional return.




Figure 3: Three price paths for a digital option that eventually matures in the money (paying \$1). amr = average modified return.




Figure 4: Averages and 95% confidence intervals (monthly modified return, return, weighted modified return and weighted return, in excess of the three-month T bill rate) for U.S IPOs in the period 1975-84, *Active issues only*. Results for 10 cumulative twelve-month periods.




Figure 5: Averages and 95% confidence intervals (monthly modified return, return, weighted modified return and weighted return, in excess of the three-month T bill rate) for U.S IPOs in the period 1975-84, *Mergers only*. Results for 10 cumulative twelve-month periods.

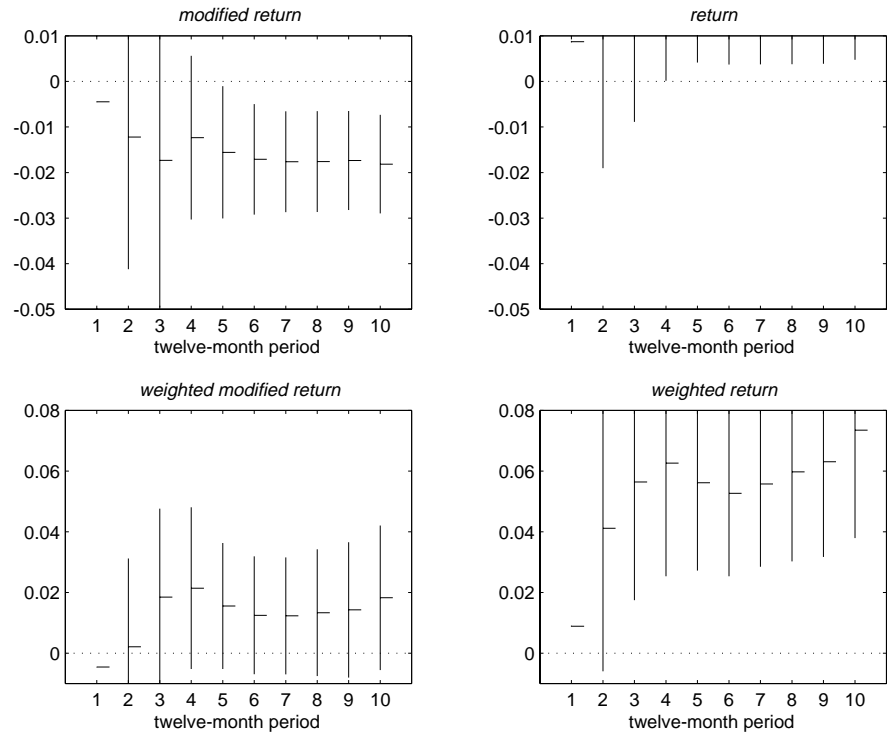


Figure 6: Averages and 95% confidence intervals (monthly modified return, return, weighted modified return and weighted return, in excess of the three-month T bill rate) for U.S IPOs in the period 1975-84, *Exchanges only*. Results for 10 cumulative twelve-month periods.




Figure 7: Average (excess) returns and 95% confidence intervals for U.S. IPOs in the period 1975-84, *losers only* (left panel: Liquidations; right panel: Unknown). Results for 10 cumulative twelve-month periods.




Figure 8: Average weighted modified (excess) returns and 95% confidence intervals for U.S IPOs in the period 1975-84, *Active issues only*. Stratified by *IPO Underpricing* decile (decile 1 = highest underpricing). Results for 10 cumulative twelve-month periods.




Figure 9: Average weighted modified (excess) returns and 95% confidence intervals for U.S IPOs in the period 1975-84, *Mergers only*. Stratified by *IPO Underpricing* decile (decile 1 = highest underpricing). Results for 10 cumulative twelve-month periods.




Figure 10: Average weighted modified (excess) returns and 95% confidence intervals for U.S IPOs in the period 1975-84, *Exchanges only*. Stratified by *IPO Underpricing* decile (decile 1 = highest underpricing). Results for 10 cumulative twelve-month periods.




Figure 11: Average weighted modified (excess) returns and 95% confidence intervals for U.S IPOs in the period 1975-84, *Active issues only*. Stratified by *industry classification* (classes are determined by first digit of Ritter's 3-digit SIC Code). Results for 10 cumulative twelve-month periods.




Figure 12: Average weighted modified (excess) returns and 95% confidence intervals for U.S IPOs in the period 1975-84, *Mergers only*. Stratified by *industry classification* (classes are determined by first digit of Ritter's 3-digit SIC Code). Results for 10 cumulative twelve-month periods.




Figure 13: Average weighted modified (excess) returns and 95% confidence intervals for U.S IPOs in the period 1975-84, *Exchanges only*. Stratified by *industry classification* (classes are determined by first digit of Ritter's 3-digit SIC Code). Results for 10 cumulative twelve-month periods.




Figure 14: Average weighted modified (excess) returns and 95% confidence intervals for U.S IPOs in the period 1975-84, *Active issues only*. Stratified by *IPO price level* decile (decile 1 = lowest IPO prices). Results for 10 cumulative twelve-month periods.




Figure 15: Average weighted modified (excess) returns and 95% confidence intervals for U.S IPOs in the period 1975-84, *Mergers only*. Stratified by *IPO price level* decile (decile 1 = lowest IPO prices). Results for 10 cumulative twelve-month periods.




Figure 16: Average weighted modified (excess) returns and 95% confidence intervals for U.S IPOs in the period 1975-84, *Exchanges only*. Stratified by *IPO price level* decile (decile 1 = lowest IPO prices). Results for 10 cumulative twelve-month periods.




Figure 17: Average weighted modified (excess) returns and 95% confidence intervals for U.S IPOs in the period 1975-84, *Active issues only*. Stratified by *entry rank in industry* (first 10 entries shown). Results for 10 cumulative twelve-month periods.




Figure 18: Average weighted modified (excess) returns and 95% confidence intervals for U.S IPOs in the period 1975-84, *Mergers only*. Stratified by *entry rank in industry* (first 10 entries shown). Results for 10 cumulative twelve-month periods.




Figure 19: Return autocorrelations (up to 6th order) for U.S IPOs in the period 1975-84. Stratified by IPO price level decile (decile 1 = lowest IPO prices). Solid lines indicate confidence bands, computed as $\pm 2/\sqrt{NT}$, where N is the number of IPOs in each decile and T the average length of the IPO aftermarket history.

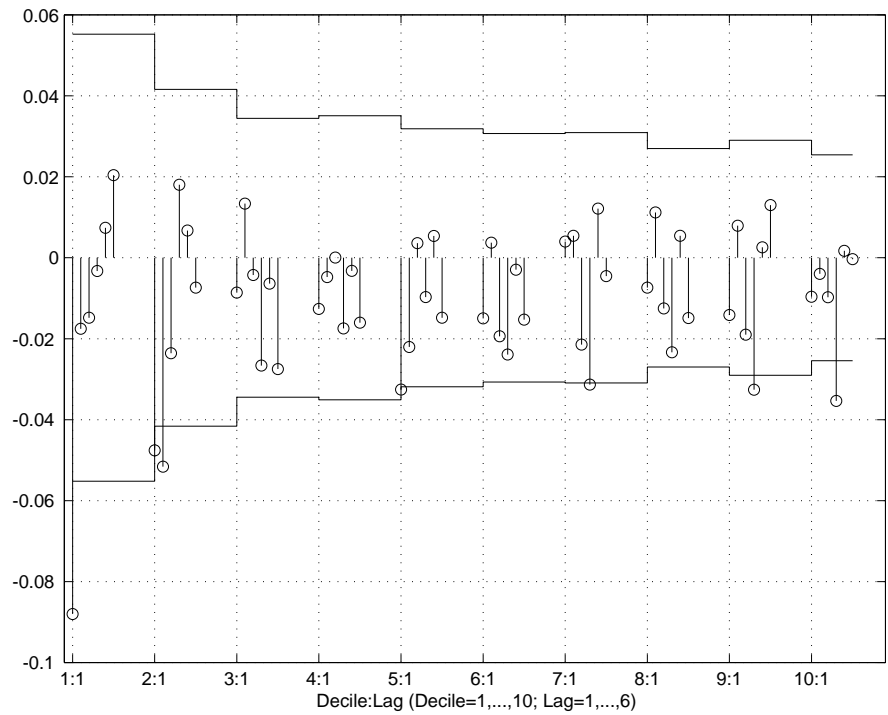


Figure 20: Return autocorrelations (up to 6th order) for U.S IPOs in the period 1975-84, *Active issues only* (as of ten years after IPO date). Stratified by IPO price level decile (decile 1 = lowest IPO prices). Results for six lags. Solid lines indicate confidence bands, computed as $\pm 2/\sqrt{NT}$, where N is the number of IPOs in each decile and T the average length of the IPO aftermarket history.




Figure 21: Return autocorrelations (up to 6th order) for U.S IPOs in the period 1975-84, *Mergers only* (as of ten years after IPO date). Stratified by IPO price level decile (decile 1 = lowest IPO prices). Results for six lags. Solid lines indicate confidence bands, computed as $\pm 2/\sqrt{NT}$, where N is the number of IPOs in each decile and T the average length of the IPO aftermarket history.




Figure 22: Simulation of a rational price path and the corresponding outcome when disturbed by noise from a Variable Reversal Delay (VRD) model, for a call option that matured in the money.




Figure 23: Notched boxplots of weighted average modified returns on winning call options. Noise from a Variable Reversal Delay (VRD) model is added incrementally from the left boxplot (zero noise) to the right one (noise scaling factor δ_n equal to 0.5). Each boxplot is based on 50 simulations of 50 price paths. Notches indicate 95% confidence intervals for the median.