

Why Don't Issuers Get Upset About Leaving Money on the Table in IPOs?

Tim Loughran

University of Notre Dame

Jay R. Ritter

University of Florida

One of the puzzles regarding initial public offerings (IPOs) is that issuers rarely get upset about leaving substantial amounts of money on the table, defined as the number of shares sold times the difference between the first-day closing market price and the offer price. The average IPO leaves \$9.1 million on the table. This number is approximately twice as large as the fees paid to investment bankers and represents a substantial indirect cost to the issuing firm. We present a prospect theory model that focuses on the covariance of the money left on the table and wealth changes. Our reasoning also provides an explanation for a second puzzling pattern: much more money is left on the table following recent market rises than after market falls. This results in an explanation of hot issue markets. We also offer a new explanation for why IPOs are underpriced.

During 1990–1998, companies going public in the United States left more than \$27 billion on the table, where the money left on the table is defined as the first-day price gain multiplied by the number of shares sold. If the shares had been sold at the closing market price rather than the offer price, the proceeds of the offering would have been higher by an amount equal to the money left on the table. Alternatively, the same proceeds could have been raised by selling fewer shares, resulting in less dilution of the preissue shareholders. The investors' profits come out of the pocket of the issuing company and its preissue shareholders.

The \$27 billion left on the table is twice as large as the \$13 billion in investment banker fees paid by the issuing companies that we study. These same companies generated profits of approximately \$8 billion in the year before going public, so the amount of money left on the table represents more than three years of aggregate profits. In some cases, the numbers are extreme:

We wish to thank Hsuan-Chi Chen and Donghang Zhang for useful research assistance. Hsuan-Chi Chen, Kent Daniel, Harry DeAngelo, Craig Dunbar, Bruce Foerster, Jason Karceski, Robert Korajczyk (the editor), Alexander Ljungqvist, M. Nimalendran, Tim Opler, Michael Ryngaert, Hersh Shefrin, William Wilhelm, Kent Womack, Li-Anne Woo, and an anonymous referee have provided useful comments, as have seminar participants at the Australasian consortium of universities videoconference, DePaul, Cornell, Georgetown, and Ohio State universities, the universities of Florida, Georgia, and Kansas, the ABN Amro Conference on IPOs in Amsterdam, the 8th Pacific Basin Accounting, Economics, and Finance Conference in Bangkok, the NYSE Equity Markets in Transition Conference in Scottsdale, the Society of Financial Studies Conference at Northwestern, the September 2000 Chicago Quantitative Alliance Conference, and the 2001 American Finance Association meetings. Address correspondence to Jay R. Ritter, P.O. Box 117168, University of Florida, Gainesville, FL 32611-7168, or e-mail: ritter@dale.cba.ufl.edu or Loughran.9@nd.edu.

In Netscape's August 1995 initial public offering (IPO) with Morgan Stanley as the lead underwriter, 5 million shares were sold to investors at \$28.00 per share. With a closing market price of \$58.25, \$151 million was left on the table. Yet in spite of this huge wealth transfer from Netscape's preissue shareholders to those lucky enough to have been allocated shares at the offer price, Netscape's major shareholders were satisfied with the pricing of the offering. (Most companies going public in the United States are relatively young firms with large blocks of equity owned by the managers. Thus throughout most of this article we will use the terms "issuing company" and "pre-issue shareholders" interchangeably.) Netscape retained Morgan Stanley as the lead underwriter for the November 1996 follow-on offering. And this reaction is not unusual. Krigman, Shaw, and Womack (2001) report that issuing firms do not view large amounts of money left on the table as an important consideration in choosing underwriters for a follow-on offering. They report that for 15 IPOs with first-day returns in excess of 60% that subsequently conducted follow-on offerings, all 15 retained the lead underwriter from the IPO.

The fact that issuers rarely complain about leaving large amounts of money on the table has long puzzled financial economists. As Brealey and Myers (1996, p. 389) state after discussing an IPO that tripled in value on its first day of trading, "Contentment at selling an article for one-third of its subsequent value is a rare quality."

Why don't issuers get upset about leaving money on the table? We propose a prospect theory answer to this question. Prospect theory assumes that issuers care about the change in their wealth rather than the level of wealth. Prospect theory predicts that, in most situations occurring in the IPO market, issuers will sum the wealth loss from leaving money on the table with the larger wealth gain on the retained shares from a price jump, producing a net increase in wealth for preissue shareholders.

Empirically we show that most of the money left on the table comes from a minority of IPOs. Indeed, although the average amount left on the table is \$9.1 million, the median is only \$2.3 million. The IPOs leaving a lot of money on the table are those where the offer price is revised upward from what had been anticipated at the time of distributing the preliminary prospectus. The offer price is increased in response to indications of strong demand, but it could have been increased even further. Thus at the same time that underpricing is diluting the preissue shareholders of these firms, these shareholders are receiving the good news that their wealth is much higher than they had anticipated. Our model is built upon this covariance of the amount of money left on the table and unanticipated wealth changes.

An example will illustrate the argument. In Netscape's IPO, James Clark, a company cofounder, held 9.34 million shares. Approximately one month before going public, Netscape filed a preliminary prospectus with the Securities and Exchange Commission (SEC). This prospectus contained a projected

number of shares to be issued and an anticipated price range for the offering. Based upon the midpoint of the file price range of \$12–\$14, the expected value of his Netscape holdings equaled \$121 million at the time that the preliminary prospectus was filed. At the closing market price on the first day of trading, his shares were worth \$544 million, a 350% increase in this component of his pretax wealth in the course of a few weeks. So at the same time that he discovered that he had been diluted more than necessary due to the large amount of money left on the table, he discovered that his wealth had increased by hundreds of millions of dollars. Would many people be upset if they found themselves in this situation?

Most times when there is a large stock price run-up, the offer price has been increased above the file price range. The empirical pattern that the first-day return is related to the revision in the offer price was first documented by Hanley (1993). IPOs where the offer price is revised upward see much higher first-day price jumps, on average, than those where the offer price is revised downward. The magnitude of the difference is large: issues where the final offer price is below the minimum of the file price range have average first-day returns of 4%, whereas those that are priced above the maximum of the file price range have average first-day returns of 32%.

The prevailing wisdom among academics for why this partial adjustment exists is based on the Benveniste and Spindt (1989) model of IPO underpricing. This dynamic information acquisition model argues that regular investors, in order to truthfully reveal their demand to an underwriter during the book-building phase of an IPO's marketing, must be rewarded with more underpricing on deals for which there is strong demand. Thus deals in which the offer price is revised upward will have greater underpricing. The Benveniste and Spindt model predicts that there should be partial adjustment of the offer price with respect to private information. It does not predict that there should be partial adjustment to public information, such as recent market movements that are readily observable to all parties. Yet, as we show, the first-day returns on IPOs are predictable based on market movements in the three weeks prior to issue. The quantitative effect is large: each 1% increase in the market during the three weeks before issue results in a first-day return that is 1.3% higher (11.3% rather than 10.0%, for example).

Because there is partial adjustment to public information, first-day returns are predictable. Following a market rise, IPOs that were in the preselling period will have higher than average expected first-day returns. Similarly when there is a market fall, IPOs that come to market in the next few weeks will have low expected first-day returns. Our prospect theory model does not distinguish between public and private information, and thus predicts that average first-day returns will be predictable based on public information. The prospect theory explanation of conditional underpricing thus leads to a theory of hot issue markets.

As many authors have documented, IPOs are underpriced, on average. We argue that this underpricing is a form of indirect compensation to underwriters. At first glance, leaving money on the table comes at a cost to underwriters. Since the percentage gross spread is typically negotiated before the final offer price is established, raising the offer price increases the revenues of underwriters. Then why do underwriters choose a lower offer price, and in so doing leave more money on the table? Investment bankers benefit in two ways. First, it makes it easier to find buyers for IPOs, reducing their marketing costs [Baron (1982)]. Second, investors will engage in rent-seeking behavior to improve their priority for being allocated shares in hot IPOs. Among other things, they do this by trading with the brokerage arm of the underwriters and overpaying for commissions.¹ This rent-seeking behavior on the part of potential IPO investors increases the revenues of the underwriter beyond that measured when focusing exclusively on the gross spread.

Although underwriters want to leave money on the table, it is unlikely that they are able to gain \$1 in pretax profits from quid pro quos for every \$1 left on the table. Thus there is some “leakage,” in that underwriters benefit more from a \$1 increase in direct fees than from leaving \$1 more on the table, even though the effect on the net proceeds received by the issuers is the same. Consequently we are trying to explain why underwriters prefer to use a relatively inefficient mechanism to collect revenue. We argue that underwriters are able to achieve higher revenue if part of their compensation is less transparent than if it all was in the form of gross spreads. This is because issuers treat the opportunity cost of leaving money on the table as less important than the direct fees [see Thaler (1980)].

Our prospect theory explanation can be recast in terms of a bargaining model in which underwriters want a lower offer price and issuing firms desire a higher offer price.² We are arguing that when unexpectedly strong demand becomes apparent during the preselling period, issuing firms acquiesce in leaving more money on the table. When demand is unexpectedly weak, issuing firms negotiate more aggressively, leaving little money on the table.

The structure of the rest of this article is as follows. Section 1 describes the data used. Section 2 describes the partial adjustment phenomenon. Section 3 provides a prospect theory explanation for why issuers aren’t upset. Section 4 presents an explanation of why underwriters intentionally underprice IPOs, on average. Section 5 addresses the question of why offer prices do not make a full adjustment to market movements in between the setting of the

¹ In a *Wall Street Journal* article on an SEC probe of mutual funds overpaying on commissions, Lucchetti (1999) states, “Other fund executives point out that higher commissions can be justified by . . . the access they can provide to initial public stock offerings.”

² The objective function of the investment banker can be viewed as benefiting from underpricing, provided it is not carried to an extreme [see Beatty and Ritter (1986) and Dunbar (2000)]. If certain conditions are satisfied, models with intertemporal maximization problems predict that some issuers may want to intentionally underprice. These models, however, do not account for dynamic price adjustment between the file price range and the final offer price.

file price range and the final offer price. Section 6 develops the implications of our model for the predictability of first-day returns and provides a theory of hot issue markets. Section 7 provides a summary and conclusions.

1. Data

Our sample is composed of 3,025 IPOs listed by Securities Data Company (SDC) from 1990 to 1998 that meet several criteria. To reduce the influence of microcap stocks, we exclude all IPOs that had a midpoint of the file price range below \$8.00 per share. Unit offerings are excluded, as are closed-end funds, real estate investment trusts (REITs), partnerships, and American depository receipts (ADRs). The first closing market price is taken from the Center for Research in Securities Prices (CRSP), as is the postissue number of shares outstanding.

Table 1 reports the number of IPOs in our sample for each year from 1990 to 1998, along with the average first-day return. The average first-day return, measured from the offer price to the closing market price, is 14% for our sample firms. Throughout the article we do not adjust for market movements in reporting first-day returns. This is because market movements are small in comparison (an average of 0.05% per day) and thus have little impact on the conclusions. The average first-day return of 14% is in line with the underpricing reported in other studies.

The last three columns of Table 1 report the proportion of IPOs priced below, within, and above the initial file price range. During our sample period approximately one-quarter are priced below, one-half within, and one-quarter above the file price range. Since our sample includes only completed deals, withdrawn offerings, most of which would have been priced below the minimum, are not reflected in the proportions.

Table 1
Number of initial public offerings, mean first-day return, and revisions from the file price range by cohort year

Cohort year	Number of IPOs	Percentage first-day return	Percentage of IPOs with OP < LO	Percentage of IPOs in the middle	Percentage of IPOs with OP > HI
1990	89	9.49	25.8	50.6	23.6
1991	250	11.32	19.6	55.6	24.8
1992	338	9.80	38.5	40.2	21.3
1993	437	11.63	21.5	54.2	24.3
1994	319	8.54	37.0	49.5	13.5
1995	366	20.37	19.7	43.7	36.6
1996	571	16.05	25.2	49.6	25.2
1997	389	13.75	30.8	45.2	23.9
1998	266	21.77	28.2	48.9	22.9
Total	3,025	14.07	27.3	48.4	24.3

The IPO information was purchased from Securities Data Company. No unit offerings, ADRs, REITs, partnerships, or closed-end funds are included in the sample. Firms with a midpoint of the initial file price range of less than \$8 are excluded from the sample. The first-day return uses the first closing price from CRSP. OP is the offer price. IPOs are categorized with respect to the minimum (LO) and maximum (HI) of the original file price range. The difference between LO and HI in the original file price range is exactly \$2.00 for 86% of the firms in our sample.

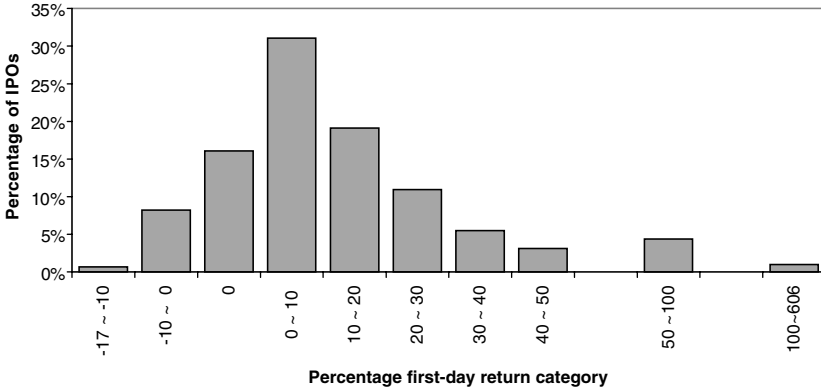


Figure 1
Histogram of first-day returns for 3,025 IPOs in 1990-1998
 Units, ADRs, REITs, closed-end funds, partnerships, and IPOs where the midpoint of the file price range was less than \$8.00 per share are excluded. The average first-day return (percentage return from offering price to first-day close) is 14.1%.

Figure 1 presents a histogram of the first-day returns. During 1990-1998, less than 1.0% of IPOs doubled in price on the first day (the proportion has been much higher after 1998). Sixteen percent of IPOs close the first day at the offer price, a feature widely attributed to stabilization activities on the part of underwriters [see Hanley, Kumar, and Seguin (1993), Jenkinson and Ljungqvist (2001), and Aggarwal (2000)].

2. The Partial Adjustment Phenomenon

Table 2 illustrates the “partial adjustment phenomenon,” first documented by Hanley (1993). In this table we categorize IPOs on the basis of the final offer price relative to the original file price range. Table 2 reports that the average IPO left \$9.1 million on the table, a number that works out to an aggregate of more than \$27 billion in 1990-1998.³ For those IPOs priced below the file range, the mean amount is \$1.5 million and the median amount is only \$0.2 million. For those priced within the file range, the mean is \$6.4 million and the median is \$2.1 million. For those priced above the file range, the mean is \$23.0 million and the median is \$12.7 million. More than 60% of

³ In all of our calculations, we do not take into account the exercise of overallotment options. The reason for this exclusion is the unreliability of this information in the SDC new issues database. Because a 15% overallotment option is present in almost all IPOs, and the overallotment option is much more likely to be exercised in full when there is strong demand, including information on the exercise of overallotment options would strengthen all of our results.

Table 2

Mean and median first-day returns, proceeds, amount of money left on the table, and revaluations from the filing to the first close, categorized by the final offer price relative to the midpoint of the file price range, for IPOs from 1990 to 1998

Item	Sample size	Mean (median) first-day return	Percentage of first-day returns that are positive	Mean (median) proceeds	Mean (median) money left on the table	Mean (median) revaluation
OP < LO	825	4.00% (0.84%)	53%	\$45.3 million (\$24.0 million)	\$1.5 million (\$0.2 million)	−\$26.4 million (−\$14.6 million)
Middle range	1,464	10.80% (7.18%)	77%	\$68.1 million (\$32.5 million)	\$6.4 million (\$2.1 million)	\$13.9 million (\$3.7 million)
OP > HI	736	31.86% (24.07%)	96%	\$93.6 million (\$46.8 million)	\$23.0 million (\$12.7 million)	\$113.4 million (\$48.2 million)
All	3,025	14.07% (7.50%)	75%	\$68.1 million (\$33.6 million)	\$9.1 million (\$2.3 million)	\$27.1 million (\$2.8 million)

The IPO information is from Securities Data Company. No unit offerings, ADRs, REITs, partnerships, or closed-end funds are included in the sample. Firms with a midpoint of the original file price range of less than \$8 are excluded from the sample. The first-day return is from the offer price to the first closing price from CRSP. OP is the offer price, and LO and HI are the minimum and maximum initial file prices. The proceeds include shares issued by both the firm and selling shareholders. The amount of money left on the table is calculated as the price change from the offer price to the closing first-day market price, multiplied by the number of shares issued (including domestic and international tranches, but excluding overallocation options). The revaluation measures the change in the wealth of preissue shareholders from the date of setting the original file price range until the close of trading on the day of going public. The revaluation is calculated as the change in price from the midpoint of the original file price range to the first closing market price for the shares retained by preissue shareholders, plus the change in price from the midpoint of the file range to the offer price for the shares sold by selling shareholders in the IPO. To calculate the number of shares retained, we use the first CRSP-listed number of shares outstanding and subtract the number of shares offered (excluding overallocation options). For 157 IPOs with multiple classes of common shares, the revaluation is computed using only the class of shares sold in the IPO. For these IPOs, the magnitude of the revaluation is underestimated. These IPOs typically have class A and class B shares, where the IPO is composed of shares with inferior voting rights. The class with superior voting rights is not traded.

the amount of money left on the table is from the 24% of IPOs priced above the file range. This skewness in the distribution of the amount of money left on the table suggests that this should be a big concern for only some issuers.⁴

The last column of Table 2 reports the revaluations of the preissue equityholders' stakes during the IPO process. This represents the change in their wealth during the interval between when the file price range is set and the close of trading on the day of issue. We compute the revaluation as the number of shares retained by preissue shareholders multiplied by the change in the value per share from the midpoint of the filing price range to the closing market price on the first day of trading, plus the shares sold by selling shareholders in the IPO multiplied by the change in the value per share from the midpoint of the file price range to the offer price.⁵ On average, this revaluation is positive,

⁴ Frequently the number of shares is revised in the same direction as the price. Logue et al. (2001, Table 1) report the average change in the number of shares offered, using the same three classifications for revisions in the offer price that we use. For IPOs where the offer price is below the minimum of the file price range, the average revision in the number of shares offered is −4%. For IPOs priced within the file range, the number of shares is revised upward by 5%, and for deals priced above the maximum, the number of shares is revised upward by an average of 9%.

⁵ As an example, Cisco Systems, a February 1990 IPO, had 9.5 million shares retained by preissue shareholders, 2.43 million newly issued shares, and 0.37 million secondary shares in its offering at \$18.00 per share. The first closing market price was \$22.25, and the midpoint of the file price range was \$14.50. So for the 9.5 million shares retained, the revaluation of \$7.75 per share resulted in a wealth gain of \$73.6 million. For the 0.37 million shares sold by existing shareholders, the revaluation of \$3.50 per share

reflecting the underpricing phenomenon. But importantly, this revaluation covaries with the amount of money left on the table. For upward revisions, where on average \$23.0 million is left on the table, this revaluation averages \$113.4 million. For downward revisions, where on average only \$1.5 million is left on the table, the revaluation is a negative \$26.4 million.

In Figure 2 we present histograms of the first-day return distributions, conditional on the revision in offer price relative to the original file price range. Inspection of these histograms shows that the distributions are very different. Conditional upon a downward revision in the offer price, the chance of a large first-day return is quite low. Conditional upon an upward revision, the probability of a high first-day return is quite high.

3. Prospect Theory and Issuer Complacency

Let's return to the example of Netscape's IPO, where cofounder James Clark held 9.34 million shares. Based on the midpoint of the file price range of \$12–\$14, the expected value of his Netscape holdings was \$121 million at the time that the preliminary prospectus was filed. At the closing market price on the first day of trading, his shares were worth \$544 million, a 350% increase in his pretax wealth in the course of a few weeks. So at the same time that he discovered that he had been diluted more than necessary due to the large amount of money left on the table, he discovered that his wealth had increased by hundreds of millions of dollars. Since he owned 28.2% of the company before going public, \$43 million of the \$151 million wealth transfer from preissue shareholders to new investors came out of his pocket. After the offering, he owned 24.5% of Netscape, but if the same proceeds had been raised by selling 2.4 million shares at \$58.25 instead of 5.0 million shares at \$28.00, he would have owned 26.3%.

Suppose instead that the offer price had been revised downward to \$6.00 per share and then jumped to \$12.50. This 108% first-day return is the same percentage increase as the actual jump from \$28.00 to \$58.25. But at a \$12.50 market price, his 9.34 million shares would be worth \$117 million, about the same wealth as he had been expecting a few weeks earlier. In this scenario, \$32.5 million would have been left on the table, of which \$9 million would have been his. Now he should be mad: he has been diluted, and there is no offsetting good news. In fact, his \$117 million ex post holding of Netscape is less than the \$121 million he had expected a few weeks earlier. We conjecture that he would be much more upset at the investment bankers for leaving \$32.5 million on the table in this scenario than he was when \$151 million was actually left on the table, but was accompanied by the good news that his wealth had increased by 350% in a matter of weeks.

resulted in a wealth gain of \$1.3 million. Thus the total revaluation was \$74.9 million for the preissue shareholders, as contrasted with the \$11.9 million left on the table.

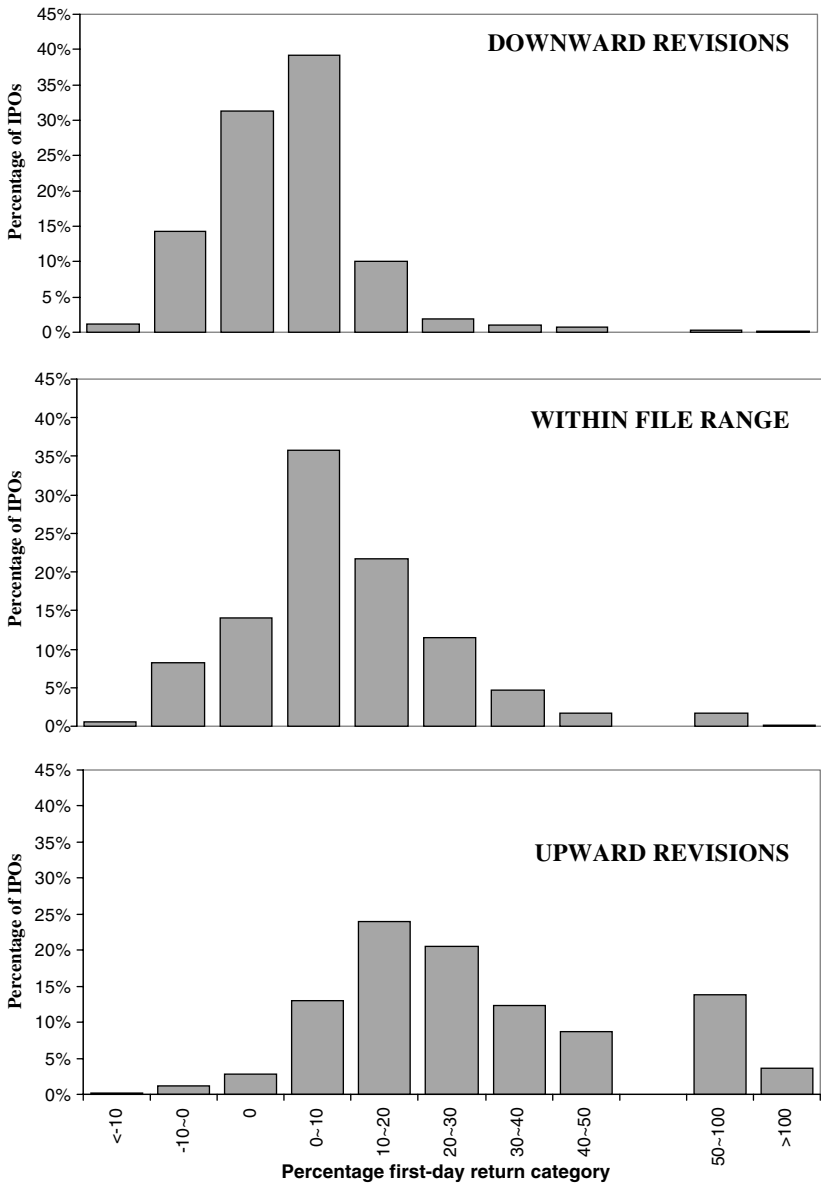


Figure 2

Conditional distributions of first-day IPO returns, 1990–1998

There are 825 IPOs with an offer price below the minimum of the original file price range, 1,464 IPOs that are priced within the range, and 736 IPOs with an offer price above the maximum of the file price range.

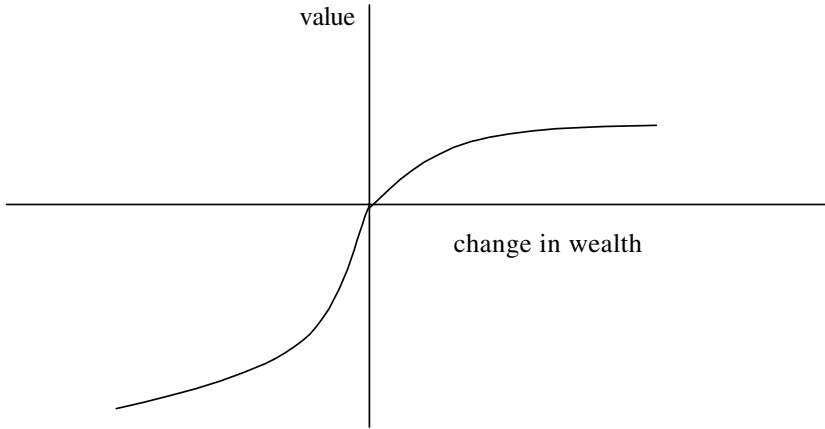


Figure 3
Prospect theory's value function, representing an individual's preferences over gains and losses relative to a reference point

The figure shows that the individual loses more value for small losses than he or she adds for a small gain of the same magnitude (loss aversion) and is risk averse (concave) for gains and risk seeking (convex) in losses.

This intuitively plausible feeling can be formalized. Prospect theory [Kahneman and Tversky (1979), Shefrin and Statman (1984)] is a descriptive theory of choice under uncertainty and is not based on normative postulates about the way people should behave. Each individual has a value function, which is similar to a utility function, but the value function is defined in terms of gains and losses rather than levels. The value function is illustrated in Figure 3. The value function is concave in gains and convex in losses, with the function being steeper for small losses than for small gains (loss aversion).

In the context of going public, gains and losses are computed relative to the price that the executives of the issuing firm have anchored on. We argue that this reference point is the midpoint of the file price range rather than the historical cost (since in many cases executives have received their shares as "sweat equity" in lieu of cash compensation).⁶

Prospect theory argues that when an individual is faced with two related outcomes, the individual can either treat them separately or as one. Figure 4 illustrates the regions for which segregation and integration of the outcomes

⁶ The June 29, 1999, IPO of e-Loan is described in the cover story of the September 6, 1999, issue of *Business Week* [Hof (1999)]. In the article, company president Janina Pawlowski boycotts the pre-IPO dinner sponsored by the lead underwriter, Goldman Sachs, because she and the other members of the management team are upset with Goldman for choosing a \$14 offer price. In the face of strong demand (the offer was oversubscribed 26 times) following the road show, the issuers had argued for a \$16 offer price at the pricing meeting. The original file price range was \$11 to \$13, so this seems to be inconsistent with our prospect theory explanation. But the key is that management anchored not on the \$12 midpoint of the file price range, but instead on \$16, the price that Softbank had been willing to pay in a proposed private equity infusion shortly before the IPO. The closing market price on the day of issue was \$37 per share, leaving \$80 million on the table.

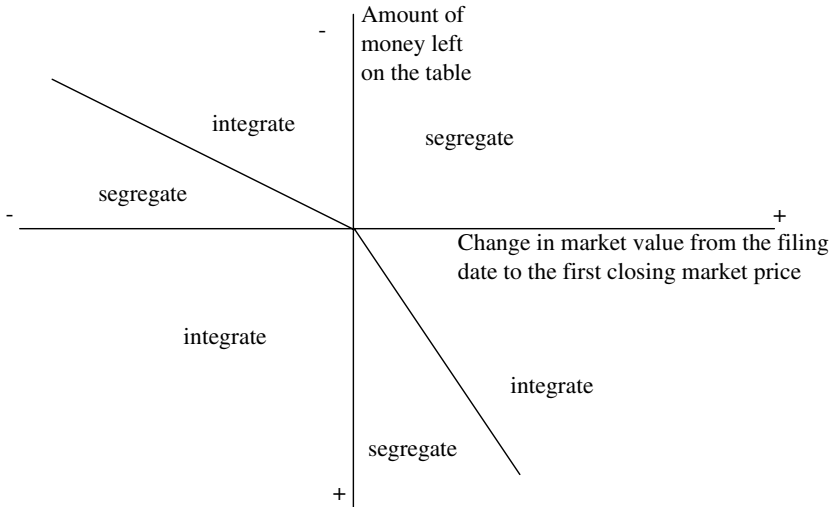


Figure 4

This figure shows the regions of integration and segregation for an individual who incurs both gains and/or losses based on revisions in valuation from the filing date to the first closing market price and due to money being left on the table. A large amount of money on the table is at the bottom of the graph, and a small (negative) amount is at the top.

will occur. Due to the concavity of the value function for gains, two gains will be segregated. This is because a person feels better about two gains than about one gain of twice the size. Because of the convexity of the value function for losses, two related losses will be integrated. But for a gain and a loss, whether the individual feels better by integrating or segregating the events depends upon their magnitudes. The lines demarcating the integrate and segregate regions of the second and fourth quadrants in Figure 4 are below the 45 degree line because of the loss aversion feature of prospect theory. By integrating the bad news (a small amount of dilution) with the good news (a high net worth increase), the stockholders can feel good about the net gain. In Figure 4, issues that are underpriced following an upward revision in the offer price will be in the lower right-hand quadrant.

In general, the wealth gain for preissue shareholder i from the revaluation is greater than his or her share of the money left on the table when the following condition is met:

$$\begin{aligned}
 &[\text{shares retained}_i + \text{secondary shares sold}_i][OP - \text{midpoint}] \\
 &+ \text{shares retained}_i[P - OP] > [P - OP][\text{secondary shares sold}_i \\
 &+ \text{primary shares sold}(\text{shares retained}_i/\text{shares retained})],
 \end{aligned}$$

where P is the market price, OP is the offer price, primary shares sold are being sold by the firm, secondary shares sold_{*i*} are existing shares being sold by shareholder i , and the shares retained without a subscript are for all shareholders combined [see Barry (1989) for related work]. If this condition is met, a preissue shareholder will find oneself in the integration region of the lower right-hand quadrant of Figure 4. For issuers with an upward revision of the offer price and underpricing, this condition will be met unless shareholder i is selling a large number of shares in the offering or the offering is huge relative to the preissue number of shares. In practice, young companies going public rarely have any secondary shares in the IPO.

We have provided a cognitive psychology argument for why some issuers will not be greatly upset with leaving money on the table in IPOs. The key element is the covariance of money left on the table and wealth gains accruing to the issuer. This is an example of the importance of framing. If issuers viewed the opportunity cost of underpricing by itself, issuers would be more resistant to severe underpricing. But because it comes as part of a package that includes the good news of an increase in wealth, there is much less resistance. We do not claim that this conditional underpricing is an optimal contract among the class of all possible contracts. Indeed, our suspicion is that book building is favored by underwriters partly because it allows them to take advantage of risk-averse issuing firms. The road show period immediately before an IPO is a high-stress period for issuing firms. The terms of the offering are subject to substantial revisions, and there is a nontrivial chance that the offering may be completely canceled due to forces outside of management's control, such as a sharp market drop. Thus there is a sense of relief with a completed offering, especially if the proceeds are higher than expected. And the media associates a large price jump with a "successful" IPO.

4. Money on the Table as Indirect Underwriter Compensation

Many reasons have been given for why underwriters underprice IPOs, on average [see Ibbotson, Sindelar, and Ritter (1994) and Jenkinson and Ljungqvist (2001) for surveys]. In general, these reasons are not mutually exclusive. Underpricing comes at a cost to the underwriters, in that higher gross proceeds from a higher offer price would result in higher fees, given that the percentage gross spread typically does not change whether the proceeds are revised upward or downward [Chen and Ritter (2000)]. We are focusing on why, conditional on strong demand becoming apparent during the selling period, underwriters intentionally underprice some IPOs severely. Underpricing does make it easier to market an IPO, but the benefit of reduced marketing efforts would seem to be minimal once an offering is oversubscribed by a factor of 10 or more.

By using underpricing as well as gross spreads, underwriters are able to get issuers to pay much higher average total costs than if all of the costs were impounded into direct fees for two reasons. The first reason is that issuers treat opportunity costs (the money left on the table) as less of a cost than the direct cost of the gross spread.⁷ The second reason is that the amount of money left on the table is state contingent. Even though underwriters may be able to capture only a fraction of the money left on the table in the form of quid pro quos, they are able to get higher total compensation than if all of their compensation was in direct fees.⁸ The necessary condition for this to be true is that the percentage quid pro quo per dollar left on the table exceeds the percentage gross spread that would be received on an extra dollar of gross proceeds.

We are not arguing that underpricing is necessary because the direct fees are insufficient to compensate underwriters for their efforts. Instead, we are arguing that with the indirect benefits of underpricing, underwriting IPOs is exceptionally profitable for investment bankers. Why, then, doesn't entry from competing investment bankers erode these profits? We believe that there are significant barriers to entry, largely due to the perceived importance of coverage by influential analysts. The importance of analyst coverage allows the high-prestige investment bankers to attract issuers in spite of leaving large amounts of money on the table.

Although opportunity costs are not viewed as equivalent to direct costs by issuers, opportunity costs do matter. In general, the smaller the offering, the less the dilution associated with a given first-day return. Thus our analysis predicts that IPOs selling a larger percentage of the firm, and with more secondary shares, should have less underpricing. Habib and Ljungqvist (2001) make a similar prediction and present supporting empirical evidence.

Lastly, it should be noted that neither the direct cost of the gross spread nor the indirect cost of money on the table shows up on an issuer's income statement. The net proceeds appear on the balance sheet, but neither the direct nor indirect costs directly affect the income statement. While most academics may find this accounting treatment irrelevant, our suspicion is that if these costs appeared on the income statement, managers would pay more attention to minimizing these costs.

⁷ Anecdotal evidence regarding the relative importance of direct costs versus opportunity costs is contained in Uttal's (1986) description of the pricing meeting for Microsoft's 1986 IPO. At the pricing meeting, the lead underwriter, Goldman Sachs, informed Microsoft that the IPO would probably start trading at \$25 or more, well above the \$16–\$19 file price range. Uttal describes the negotiation of the offer price in increments of a dollar per share, and of the gross spread in increments of a penny per share: Microsoft and Goldman Sachs "had no trouble agreeing on a final price of \$21. . . . Having agreed fairly easily over dollars, the two sides bogged down over pennies." Microsoft and its selling shareholders paid \$3.66 million in direct fees and left \$18.87 million on the table, with the stock price closing at \$27.75 on the first day of trading.

⁸ Habib and Ljungqvist (2001), however, present empirical evidence that the opportunity cost of underpricing has a one-to-one marginal trade-off with direct expenses.

Table 3
IPOs categorized by prior market returns, 1990–1998

Item	Sample size	Mean (median) first-day return	Mean (median) money left on the table	Mean prior 15-day value weighted index return
Prior < 0.0%	948	10.02% (4.36%) (SD = 17.0%)	\$5.6 million (\$1.3 million) (SD = \$13.3 million)	-1.84%
Middle range	1,003	13.16% (7.30%) (SD = 18.7%)	\$8.5 million (\$2.2 million) (SD = \$21.7 million)	0.99%
Prior > 2.0%	1,074	18.49% (11.11%) (SD = 31.3%)	\$12.7 million (\$4.1 million) (SD = \$32.1 million)	3.68%
All	3,025	14.07% (7.50%) (SD = 23.8%)	\$9.1 million (\$2.3 million) (SD = \$24.2 million)	1.06%

No unit offerings, ADRs, REITs, partnerships, or closed-end funds are included in the sample of 3,025 IPOs from 1990 to 1998. Firms with a midpoint of the file price range of less than \$8 are excluded from the sample. The first-day return uses the first closing price from CRSP. OP is the offer price. Prior returns are calculated over the 15 trading days before the offer date. The three categories are formed on the basis of the prior value weighted CRSP NYSE/AMEX/Nasdaq market return (15 trading days before the offering). The cutoffs of 0 and 2% are the integer values that divide the sample into three roughly equal-size subsamples. If Nasdaq index returns are used instead, the prior 15-day mean returns are virtually identical. SD is the standard deviation.

5. Why Is There Partial Adjustment to Market-Wide Information?

5.1 Empirical patterns

In Table 3 we categorize IPOs by the market movement in the three weeks (15 trading days) prior to issue.⁹ We report that following market declines, the average first-day return is 10.0%, whereas after periods when the market has risen by at least 2.0%, the average first-day return is 18.5%. The ability of recent market movements to predict first-day returns has been documented for decades [Logue (1973, Table 1), Hanley (1993, Table 3), Benveniste, Ljungqvist, Wilhelm, and Yu (2001)].¹⁰ Thus underwriters do not fully adjust the offer price with respect to public information. Indeed, the evidence suggests that there is very little adjustment of the offer price with respect to market movements. The difference in average value-weighted market returns between our top and bottom categories in Table 3 is 5.52%. The difference

⁹ During January 1997–August 1999, the length of time between filing a registration statement (form S-1 for most IPOs) with the SEC and going public has averaged 78 days, for IPOs where this length of time took less than 180 calendar days (taking more than six months indicates that the original plans were postponed at some point). Typically four to five weeks before the expected offer date, the issuer files an amended registration statement S-1/A giving a file price range. This preliminary prospectus (also known as a red herring, due to the red letters on the cover page identifying it as preliminary) is then distributed to potential investors. The firm may file further amendments as well. The SEC mandates that the final offer price must be within 20% below the minimum to 20% above the maximum of the latest file price range. Our use of 15 trading days (three weeks) is a minimal measure of the length of time between the setting of the original file price range and the offer date.

¹⁰ Maksimovic and Unal (1993), however, do not find predictability when examining the first-day returns on the IPOs of mutual financial institutions converting to stock ownership. The offer price for these IPOs must be approved by regulators.

in average first-day returns is 8.47%. Following market rises, issuers leave more than twice as much money on the table as following market declines (\$12.7 million versus \$5.6 million).

In Figure 5 we sort IPOs by both the revision in the offer price from the file price range and by prior market movements. Inspection of Figure 5 discloses that, conditional on revisions, lagged market returns predict first-day returns.

While we have shown in Table 2 that first-day returns are predictable using the revision in the offer price from the file price range, these revisions are themselves somewhat predictable, as illustrated in Figure 6. Here, for each month between March 1991 and August 1998, the percentage of IPOs that are priced above their file price range is graphed on the top and the percentage of IPOs that are priced below their file price range is graphed on the bottom. (The sample starts in March 1991 because many of the months from January 1990 to February 1991 had few IPOs, as was the case in September and

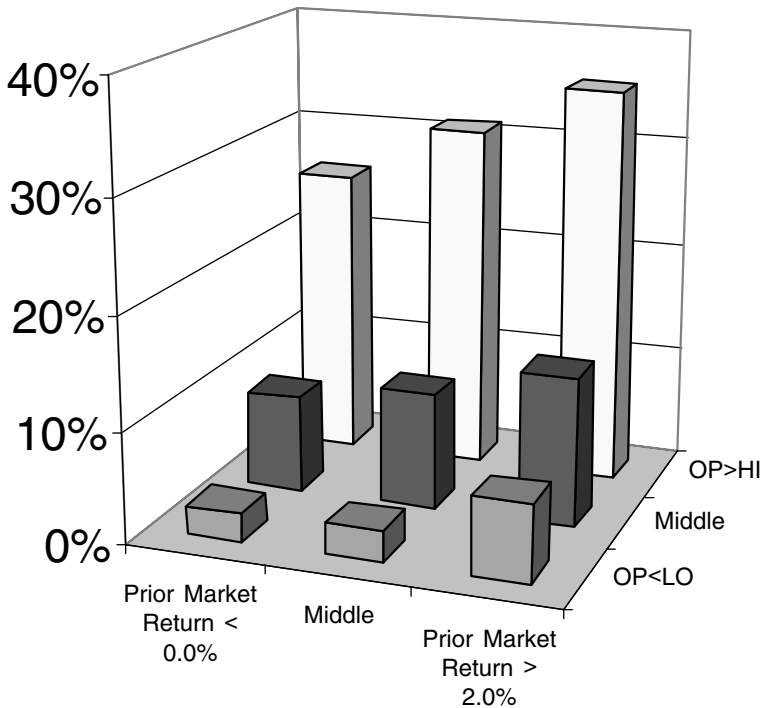


Figure 5

First-day returns categorized by prior market movements and revision in the offer price

The prior market movement is measured using the CRSP value-weighted market index in the 15 trading days prior to the IPO. Revisions in the offer price are measured relative to the original file price range, with the middle category including all IPOs that have an offer price within the original file price range. The sample is composed of 3,025 IPOs from 1990 to 1998 meeting the following requirements: no units, ADRs, REITs, closed-end funds, or IPOs with a midpoint of the original file price range of less than \$8.00. The sample size for each category is given in Table 4.

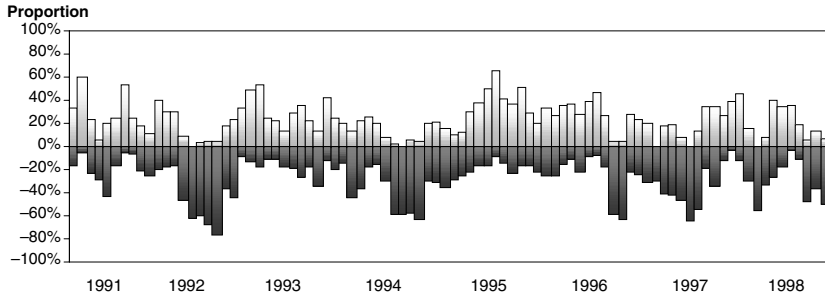


Figure 6
Percentage of IPOs, per month, with an offer price above (top) or below (bottom) the file price range
 The sample is composed of the 2,903 IPOs between March 1991 and August 1998. January 1990–February 1991 and September–December 1998 are excluded because many of the months had less than five IPOs, making the monthly percentages very sensitive to individual deals. All months from March 1991 to August 1998 had at least 11 IPOs meeting our sample selection criteria, except for January 1995, which had eight IPOs.

October 1998.) For a given month, these two percentages will generally add up to less than 100% because of other offerings that are priced within the file price range. Inspection of the figure discloses that there is persistence in the proportion of upward and downward revisions from month to month. Indeed, we can compute the average percentage change from the midpoint of the file price range to the final offer price each month. The first-order autocorrelation of this series of monthly observations is 0.61 ($p = 0.0001$) and the autocorrelation at two lags is 0.28 ($p = 0.0075$).

Information about previous first-day returns, which of course is public information, can be used to predict the average first-day return in a month. In Figure 7 we show the average first-day return by month for the same period

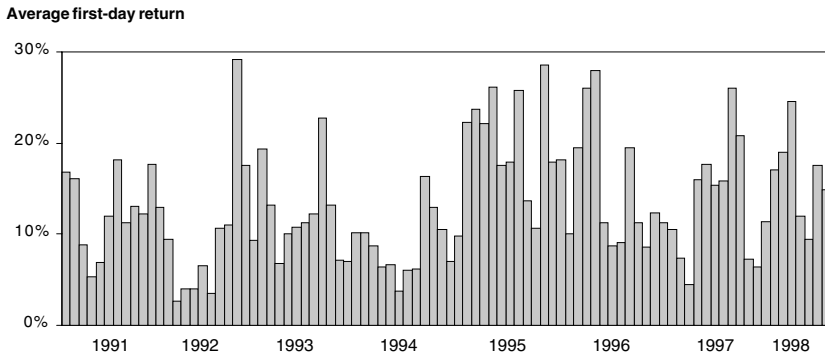


Figure 7
The equally weighted average first-day return per month from March 1991 to August 1998 for the 2,903 IPOs meeting our sample selection criteria
 A total of 122 IPOs are dropped because of our shorter sample period. The first-order autocorrelation coefficient is 0.50, with a p -value of 0.0001.

as in Figure 6. For the March 1991–August 1998 period, the first-order autocorrelation of monthly average first-day returns is 0.50 ($p = 0.0001$) and the second-order autocorrelation is 0.18 ($p = 0.089$). This persistence has been previously documented in the literature [Ibbotson and Jaffe (1975), Ibbotson, Sindelar, and Ritter (1994)]. Not surprisingly, there is contemporaneous correlation of monthly average first-day returns and the average revision in the offer price during a month. The contemporaneous correlation coefficient is 0.77, which is significantly different from zero with a p -value of 0.0001.

5.2 The dynamic information acquisition and prospect theory predictions

There are a number of alternative explanations for why underwriters do not fully adjust the offer price to information about the state of demand. These theories are not mutually exclusive. The Benveniste and Spindt (1989) dynamic information acquisition hypothesis predicts that regular investors should be rewarded for revealing their private information. In the dynamic information acquisition framework, conditional underpricing is a result of an incentive compatibility constraint. There should, however, be full adjustment to public information, such as whether the market went up or down during the weeks between when the file price range was set and the final offer price was set.

Our prospect theory explanation of the partial adjustment phenomenon does not distinguish between private and public information, so the partial adjustment of the offer price with respect to prior market movements is as predicted. Investment bankers can selectively underprice some IPOs by combining the bad news that there has been excessive dilution with the good news that they are wealthier than expected.

5.3 Optimal risk sharing

Any theory (including our prospect theory explanation) that has partial adjustment to public information implies that the issuer is being partly insured against market movements. It should be noted that publicly traded instruments, such as put options on the market and various industry indices, exist and they would appear to be the most efficient manner for providing this insurance. Thus any complete theory must include an explanation for why these instruments aren't used to insure the issuer against market movements.

A simple extension of Mandelker and Raviv (1977) would show that with a risk-averse principal (the issuer) and a risk-neutral agent (the underwriter), and with no private information, the optimal contract is for the issuer to be given fixed net proceeds. The underwriter would bear the risk of changes in valuation during the preselling period. With private information about firm value that may leak out during the preselling period, the analysis would have to be modified. With the issuing firm possessing private information, the net proceeds to the issuing firm should not change based upon public information

about market movements, but there should be partial adjustment with respect to firm-specific valuation changes [Baron (1982)].

A more realistic model must take into account that there are three parties, not two, involved. Changes in valuation must be apportioned among issuers, underwriters, and investors. How changes in valuation are borne between underwriters and investors determines the adjustments in gross spreads versus adjustments in the amount of money left on the table.

Both our prospect theory explanation and the optimal risk-bearing explanation predict that there will be state-contingent wealth transfers from issuers to the two other parties involved, underwriters and investors. The optimal risk-sharing explanation cannot explain why publicly traded instruments, such as put options on a market index, are not used to insure issuers against adverse market movements. Our prospect theory explanation, on the other hand, argues that issuers let down their guard when there is good news and do not bargain aggressively for a higher offer price, whether this is due to public information or not.¹¹

In addition to making predictions about the level of wealth transfers in various states of the world, the optimal risk-sharing and prospect theory explanations also make predictions about the form of payments (spreads versus money on the table). The optimal risk-sharing explanation for leaving money on the table must explain why investors rather than underwriters perform the risk-bearing function. If underwriters bore the risk, then percentage spreads would be higher in the good state of the world than in the bad state of the world. Without making additional auxiliary assumptions, we do not have an obvious explanation for why this does not occur if optimal risk sharing is the salient factor for explaining why the amount of money left on the table varies in a predictable manner.

Our prospect theory explanation states that issuers make a distinction between direct costs (spreads) and opportunity costs (money on the table). If underwriters increased their percentage spreads in the good state of the world, issuers would bargain harder than if the wealth transfer was in the form of an opportunity cost. Thus underwriters find that leaving money on the table may be more advantageous, in that rent-seeking behavior by investors will allow underwriters to recapture this wealth transfer from issuers.

5.4 Alternative explanations for the partial adjustment phenomenon

This article presents an agency theory explanation for the underpricing of IPOs. Previous agency theories of investment bankers taking advantage of issuers [i.e., Baron (1982)] have investment bankers using their superior knowledge about private information to underprice IPOs. As with Benveniste

¹¹ Sternberg (1989) contains a bargaining model to rationalize the partial adjustment phenomena, but he assumes a Rubenstein split-the-difference equilibrium.

and Spindt, Baron's model doesn't make the prediction that there will be partial adjustment to public information.

An explanation offered by investment bankers as to why the offer price adjusts only partially is as follows. Potential investors anchor on the midpoint of the file price range, just as issuing firms do. If the offer price is raised too far above this, some potential investors will defect [see Hof (1999) for an example of this argument], even if the increase is in response to publicly available information. Thus, because of the anchoring, underwriters are loath to adjust the offer price too much. There are two variants of this argument, which we call the unsophisticated and the sophisticated variant.

The unsophisticated version is just stating that a higher offer price deters some buyers. In other words, there is a negatively sloped demand curve. Raising the offer price reduces the excess demand. Although the underwriter can exercise more discretion in allocating underpriced shares if there is excess demand, it is not clear why this is of great advantage to the issuer.

The sophisticated version of the argument that raising the offer price in the face of strong demand induces some buyers to defect is based on Welch's (1992) cascades argument. The logic is that if investors are paying attention to what other investors are doing, increasing the offer price is risky because if some investors decide not to buy, many others may suddenly withdraw their purchase orders as well. We are skeptical of this argument for the following reason. Institutional investors are well aware of the partial-adjustment phenomenon, so that an increase in the offer price generally results in increased demand because it is signaling to investors that other investors want to buy the IPO. Because of this, once a file price range has been set there tends to be a *positively* sloped demand curve. Two axioms among IPO investors are "Cut the deal, cancel my order" and "Increase the deal, double my order" [Fitzgibbon (1998)].

One reason that is sometimes given by academics for why IPOs are underpriced, on average, is that high-quality firms are trying to signal their type by intentionally leaving money on the table. While this may partly explain why some IPOs are underpriced, it does not explain why there is partial adjustment to public, as opposed to private information.

An alternative explanation for partial adjustment is sometimes voiced by underwriters. This is the "leaning against the wind" hypothesis. Investors in the IPO market tend to overreact. Thus when the market is too receptive, the market price is bid up too high, above its long-run value. Underwriters price the issues with long-run value in mind. This leaning against the wind theory has the testable implication that there should be a negative correlation of first-day returns and subsequent long-run returns.

This implication is tested in Table 4, where we report three-year buy-and-hold returns for the IPOs in our sample after categorizing them on the basis of revisions in the offer price, and prior market movements. We also report the compounded value-weighted market returns over the same holding periods,

Table 4
Number of observations, three-year buy-and-hold returns, and wealth relatives for IPOs from 1990 to 1998, categorized by the final offer price relative to the file price range and prior 15-day market return

	Offer price relative to file range			All
	OP < LO	Middle range	OP > HI	
Prior 15-day market return				
Return < 0.0%	278	499	171	948
	59.5%	38.7%	61.5%	48.9%
	63.0%	58.6%	58.0%	59.8%
	0.98	0.87	1.02	0.93
Middle range	284	480	239	1,003
	62.1%	65.7%	70.1%	65.7%
	60.9%	62.6%	61.7%	61.9%
	1.01	1.02	1.05	1.02
Return > 2.0%	263	485	326	1,074
	38.4%	31.7%	71.4%	45.4%
	63.4%	64.8%	66.6%	65.0%
	0.85	0.80	1.03	0.88
All	825	1,464	736	
	53.7%	45.2%	68.7%	
	62.4%	62.0%	63.0%	
	0.95	0.90	1.03	

The sample consists of 3,025 IPOs from 1990 to 1998. No unit offerings, ADRs, REITs, partnerships, or closed-end funds are included in the sample. Firms with a midpoint of the initial file price range of less than \$8 are excluded from the sample. OP is the offer price, and LO and HI are the minimum and maximum file prices in the initial filing. Prior returns are calculated over the 15 trading days before the offer date. The three prior return categories are formed on the basis of the prior value-weighted NYSE/AMEX/Nasdaq market return. For each category, four numbers are reported. The top number is the sample size. The second number is the mean three-year buy-and-hold return on the IPOs. Three-year buy-and-hold returns are measured from the first CRSP-listed closing market to the earlier of the three-year anniversary, the delisting date, or December 31, 1999. The third number is the average compounded CRSP value-weighted market index return over the same three-year intervals. The bottom number in each cell is the wealth relative, calculated as the average gross return on IPOs divided by the average gross market return.

and the wealth relatives, computed as the average gross return on the IPOs divided by the average gross market return. Inspection of the table discloses no obvious patterns, with most of the wealth relatives near one, whether the offer price had been adjusted up or down.¹² The lack of a pattern is also present when, in unreported results, we use style benchmarks (size and book-to-market). The lack of any reliable patterns is inconsistent with the leaning against the wind hypothesis.¹³

Other explanations that probably have some merit include the following: If an issuer questions an underwriter about severe underpricing, the investment

¹² If we categorize IPOs by a finer partitioning of the revision of the offer price from the file range, the lack of patterns in long-run returns continues to hold.

¹³ Evidence supporting the leaning against the wind hypothesis is contained in Ritter (1991) and Krigman, Shaw, and Womack (1999). Ritter reports that IPOs with the highest first-day run-ups have the lowest long-run size-adjusted returns. Part of the reason is that Ritter's results, using IPOs from 1975 to 1984, are dominated by microcap IPOs. Ritter includes all IPOs with an offer price of \$1.00 or more, whereas we exclude IPOs with a midpoint of the file price range less than \$8.00. Krigman, Shaw, and Womack report that IPOs with first-day returns in excess of 60% subsequently underperform, whereas offerings with moderate first-day returns outperform a size-matched benchmark during their first year. Their subsample of IPOs with big first-day returns from January 1988 to May 1995 contains only 33 IPOs. The late 1990s saw a number of firms with high first-day returns (such as Yahoo!) produce extremely high returns during the next few years. Consistent with our evidence, Logue et al. (2001) also report no long-run predictive ability when IPOs are categorized by revisions in the offer price.

banker will always be willing to argue that the price jump was due to a successful job of marketing the issue by the investment banker. And issuers are well aware that the price of stocks that have jumped a lot in price can come down by the time that a lockup provision expires. There is also an element of investment banker psychology that comes into play. When there is weak demand for an offering, investment bankers do not like having to deliver bad news to an issuer. To minimize the confrontation, they do not lower the offer price (and the proceeds) as much as would be needed to maintain a target of 10–15% underpricing. If the offer price is lowered too much, there is also the danger that the issuing firm will decide to cancel the offering and seek alternative financing.

6. A Theory of Hot Issue Markets

One of the puzzles regarding IPOs is the existence of “hot issue” markets. Following Ibbotson and Jaffe (1975), a hot issue market is defined as a month in which the average first-day return is above the median month’s average first-day return.¹⁴ There is strong positive serial correlation in the monthly average first-day returns. Figure 7 shows the monthly average first-day returns on IPOs. The first-order autocorrelation coefficient is 0.50 and the second-order autocorrelation coefficient is 0.18.

Currently the literature offers no explanation that is consistent with rational behavior on the part of investors that can generate this positive autocorrelation. When there is a rise in the market, our prospect theory explanation predicts an increase in the expected underpricing of all IPOs that are in the selling period, whether they will be going public tomorrow or five weeks later. Thus we can explain the autocorrelation of average first-day returns. It is an equilibrium explanation in that no one is acting irrationally.

Private information is serially uncorrelated, so there should be no autocorrelation in first-day returns according to the dynamic information acquisition model. Prospect theory, on the other hand, does not distinguish between private and public information as a reason for wealth changes. Because selling periods overlap, there will be autocorrelation in first-day returns. First-day returns will be higher following market rises, and this effect will be present in all IPOs where the selling period includes the period of the market rise. Thus market rises will be followed by high average first-day returns for one to two months, since the time between choosing a price range on which issuers anchor and actually going public is typically in this range.

These ideas can be formalized in Equations (1) and (2). Equation (1) expresses the offer price as equal to the midpoint of the file price range plus several additional terms reflecting new information that arrives dur-

¹⁴ Other authors implicitly or explicitly use some measure of volume. Typically the volume of IPOs lags average first-day returns by several months [Ibbotson, Sindelar, and Ritter (1994)].

ing the book-building period. This book-building period starts at time 0 and ends at time T, the date of setting the final offer price. These additional terms reflect the market return, $r_{0,T}^{\text{market}}$, additional public information such as industry-specific movements, $r_{0,T}^{\text{public}}$, and firm-specific idiosyncratic information, $r_{0,T}^{\text{private}}$. By definition, this last component is uncorrelated across firms, no matter when they go public. The market return and other public-information return for date t , however, affect all IPOs whose book-building period includes time t .

$$OP_i = \text{midpoint}_i [1 + \alpha_1 r_{0,T}^{\text{market}} + \alpha_2 r_{0,T}^{\text{public}} + \alpha_3 r_{0,T}^{\text{private}}] + \tau_i \quad (1)$$

$$P_{\text{market},i} = 1.14 \cdot \text{midpoint}_i [1 + \beta r_{0,T}^{\text{market}} + \gamma_2 r_{0,T}^{\text{public}} + \gamma_3 r_{0,T}^{\text{private}}] + \varepsilon_i \quad (2)$$

Equation (2) expresses the market price as a function of the original valuation multiplied by additional terms reflecting new information that arrives during the book-building period. In Equation (2), the scalar of 1.14 is present because the midpoint of the file price range is a biased estimator of the first closing market price; our sample is underpriced by 14%, on average.

The predictions of the Benveniste and Spindt dynamic information acquisition model (hereafter, B-S) and our prospect theory model (hereafter, P-T) for the coefficients of Equations (1) and (2) are as follows: in Equation (1), B-S predicts that α_1 should equal the Sharpe–Lintner beta, α_2 should equal the relevant factor sensitivity, and α_3 should be between 0 and 1 (partial adjustment to private information). P-T predicts that α_1 should be a fraction of beta, α_2 should be a fraction of the factor sensitivity, and α_3 should be between 0 and 1, since there should be partial adjustment to both private and public information. We can distinguish between the two theories based on the difference in the coefficients in Equation (2) versus Equation (1). If B-S is a full description, then the coefficient on α_1 in Equation (1) should equal that on β in Equation (2). If P-T is descriptive of the process, then the coefficient on α_1 in equation (1) should be closer to zero than β in Equation (2).

Before testing these predictions, both equations are divided by the midpoint of the file price range so that the dependent variables are expressed as revisions from the midpoint. [The transformations are given as equations (1') and (2') in the appendix.] As a result of this division, the explanatory variables are expressed as returns rather than returns times the midpoint. Since the only observable variable that we use is the market return, we run univariate regressions. The residual term in these regressions includes not only the error term, but also the terms representing other public information and private information. In other words, we run single-variable regressions with the market return as the explanatory variable. These are reported in Table 5, after we subtract 1 from each side and multiply by 100 to convert into percentage changes.

When using the “return” on the IPO from the midpoint of the file price range to the first closing market price, the coefficient on the market return

Table 5

OLS regressions with the percentage market return in the 15 trading days before the IPO as the explanatory variable: $R_i = a_0 + a_1 R_m + e_i$

Dependent variable	a_0	a_1	R^2_{adjusted}	N
Panel A: All months				
$\left[\frac{\text{Close} - \text{OP}}{\text{OP}} \right] \times 100$	12.66 (33.91)	1.33 (6.60)	0.02	3,025
$\left[\frac{\text{OP} - \text{midpoint}}{\text{midpoint}} \right] \times 100$	-1.47 (-3.98)	0.76 (5.56)	0.01	3,025
$\left[\frac{\text{Close} - \text{midpoint}}{\text{midpoint}} \right] \times 100$	13.09 (17.41)	2.37 (7.47)	0.02	3,025
Panel B: IPOs following negative 15 trading day market returns				
$\left[\frac{\text{Close} - \text{OP}}{\text{OP}} \right] \times 100$	10.47 (11.37)	0.25 (0.63)	0.00	948
$\left[\frac{\text{OP} - \text{midpoint}}{\text{midpoint}} \right] \times 100$	-0.01 (-0.01)	1.51 (4.58)	0.02	948
$\left[\frac{\text{Close} - \text{midpoint}}{\text{midpoint}} \right] \times 100$	12.22 (6.51)	2.03 (2.73)	0.01	948
Panel C: IPOs following positive 15 trading day market returns				
$\left[\frac{\text{Close} - \text{OP}}{\text{OP}} \right] \times 100$	11.97 (14.03)	1.66 (3.98)	0.01	2,077
$\left[\frac{\text{OP} - \text{midpoint}}{\text{midpoint}} \right] \times 100$	-0.88 (-1.14)	0.50 (1.80)	0.00	2,077
$\left[\frac{\text{Close} - \text{midpoint}}{\text{midpoint}} \right] \times 100$	13.21 (7.96)	2.36 (3.60)	0.01	2,077

Initial public offerings from 1990 to 1998 meeting the criteria described in the previous tables are used in OLS regressions with the explanatory variable being the compounded percentage return on the CRSP value-weighted index during the 15 trading days prior to the offer date. In each panel the dependent variable in the top regression is the percentage return on the first day of trading. In the middle regression the dependent variable is the revision in the offer price relative to the midpoint of the initial file price range, as recorded by Securities Data Company. The dependent variable in the bottom regression is the percentage change from the midpoint to the closing market price on the first day of trading, as recorded by CRSP. T -statistics are calculated using White's (1980) method. In panel A, the Wald test statistic for the hypothesis that the slope coefficients of 0.76 and 2.37 are identical is 63.47 ($p = 0.0001$) in a seemingly unrelated regression framework. The correlation of the residuals from the second and third regressions in panel A is 0.837.

during the 15 trading days prior to the offer date implies a beta of 2.37 (see row 3 of panel A). While high, this is consistent with the findings of Clarkson and Thompson (1990) and Chan and Lakonishok (1992) that the beta of IPOs shortly after going public is in excess of 2. Using the percentage revision in the offer price as the dependent variable, the slope coefficient is only 0.76 (row 2 of panel A). The ratio of coefficients, 0.76/2.37, is only 0.32, far below the 1.00 that would be implied by the B-S model.¹⁵ This implies

¹⁵ Since some of the IPOs where downward revisions occur are withdrawn, some of the population of IPOs with negative error terms are deleted from the sample. The resulting ordinary least squares (OLS) slope coefficients will be biased toward zero. It is not obvious whether the 0.32 ratio is biased upward or downward, however.

that only one-third of the public information about market returns during the book-building period is incorporated into the offer price. This is consistent with our prospect theory explanation being economically important relative to the dynamic information acquisition model.

In row 2 of Table 5's panel A, where the dependent variable is the percentage revision of the offer price from the midpoint of the file price range, the intercept can be interpreted as the average percentage change in the offer price relative to the midpoint in a flat market. The intercept of -1.47 implies that, in a flat market, the offer price is only 1.47% below the midpoint. Thus the midpoint is quite close to being an unbiased estimate of the final offer price, conditional upon the IPO being completed, since withdrawn deals are not in our sample. In panels B and C, where we categorize IPOs on the basis of whether the market had gone down or up prior to the offer, the intercepts in the offer price revision regressions are even closer to zero. Furthermore, in all three panels, the intercepts in the regressions using the closing market price in the dependent variable are close to the 14% unconditional average return for our sample.

What autocorrelation of monthly average first-day returns does this imply? The answer depends on the relative magnitude of the variance of the error terms, the variance of public information, and the variance of private information. To take extreme cases, if all information is private, the autocorrelation will be zero. If all information is public information about market returns and the partial adjustment coefficient is zero, we have a situation analogous to autocorrelation of an index introduced by nonsynchronous trading [Atchison, Butler, and Simonds (1987)].

In the appendix we derive the predicted autocorrelation of the monthly average first-day returns. Obviously the exact numbers depend on parameter values. Assuming that the lagged market return is the only publicly available information, we derive a first-order autocorrelation coefficient in the vicinity of 0.1 to 0.2. Although positive, these values are noticeably below the empirical coefficient of 0.50. In our analysis we do not include publicly available information about industry returns. The extreme underpricing of many internet IPOs in 1999, at the same time that most non-internet IPOs were not severely underpriced, indicates that industry factors are economically important. Inclusion of lagged industry returns would undoubtedly increase the predicted autocorrelation.

Lowry and Schwert (2002) document that there is much more of an adjustment downward following market declines than upward following market increases. In panels B and C of Table 5, we confirm this. Panel B reports regression results for IPOs following market declines. The coefficient of 1.51 (t -statistic of 4.58) on the market return in row 2 shows that when the market declines by 1%, the offer price is cut by 1.51%, on average. Panel C reports regression results following market increases. The coefficient of 0.50 (t -statistic of 1.80) on the market return in row 2 shows that the offer price

is increased only minimally following market rises. Instead, the row 1 coefficient of 1.66 (*t*-statistic of 3.98) shows that investors capture most of the increased valuation following the market rise.

Panels B and C of Table 5 demonstrate that there is an asymmetric response of the offer price to market movements during the road show period. When the market declines, the issuing firm bears a reduction in the proceeds. When the market increases, investors reap a windfall. This asymmetric response is inconsistent with optimal risk sharing (where a symmetric response would be predicted), but consistent with our prospect theory explanation. Because of the increase in wealth that issuers will receive due to unanticipated market increases, issuers do not bargain hard for an offer price increase when the market goes up. Underwriters take advantage of this weak bargaining effort by not increasing the offer price. Instead, investors capture higher returns that are predictable based on prior market movements.

7. Conclusion

This article provides an explanation for why issuers don't object to large amounts of money being left on the table in IPOs. The average IPO leaves \$9.1 million on the table, an amount equal to years of operating profits for many of the companies going public, and an amount equal to approximately twice the direct fees paid to underwriters. The explanation offered here involves several parts. Most IPOs leave relatively little money on the table. The IPOs where a lot of money is left on the table are generally those where the offer price and market price are higher than had originally been anticipated. Thus the minority of issuers losing wealth via leaving large amounts of money on the table are generally simultaneously discovering they are wealthier than they expected to be. By integrating the loss with the gain, they are left happy, even though they have just been victimized. Our explanation emphasizes the covariance of the money left on the table and changes in the wealth of the issuing firm's decision makers.

This model also explains a second empirical pattern. Because offer prices only adjust partially to public information, first-day returns are predictable based on lagged market returns. Because the lagged market returns are correlated for IPOs whose preselling periods overlap, this generates autocorrelation in the first-day returns. Thus we have developed an equilibrium theory of hot issue markets. Our analysis predicts that hot and cold issue periods will be present as long as book building is used. For related evidence in a non-U.S. context, see Derrien and Womack (2002), who report that IPO auctions have less state-contingent underpricing than other mechanisms used in France.

We also offer an explanation for the IPO underpricing phenomenon. We argue that leaving money on the table is an indirect form of underwriter compensation, because investors are willing to offer *quid pro quos* to underwriters to gain favorable allocations on hot deals. Underpricing is an indirect cost to

issuers, however, and they acquiesce in severe underpricing only when they are simultaneously getting good news in the form of unanticipated wealth increases. In general, because issuers do not treat the opportunity cost of underpricing as equivalent to the same direct costs, underwriters are able to achieve higher total compensation than if all of the costs borne by issuers were bundled as direct fees.

Our analysis can be viewed as a theory of bargaining in IPOs. Our prospect theory analysis provides an explanation for why issuers bargain hard over the offer price in a bad state of the world, whereas they are pushovers in bargaining over the offer price in a good state of the world. This differential bargaining effort results in much higher underpricing when demand for an IPO is strong, whether or not this is due to publicly observable market returns during the road show period.

While we believe that our prospect theory framework provides an explanation for some of the puzzling patterns with IPOs, we do not claim to be able to explain everything. Our framework should be viewed as complementary to alternative frameworks rather than mutually exclusive. In particular, we would expect that venture capitalists, with their years of experience at taking firms public, might be less susceptible to psychological factors affecting their aggressiveness in bargaining. Yet we are unaware of any evidence that venture capital-backed IPOs are less subject to the psychological effects that we discuss than are other IPOs. Furthermore, we do not explain why issuers don't avoid investment bankers with a reputation for leaving large amounts of money on the table in prior deals.

Lastly, 1999 saw an unprecedented amount of money left on the table: \$37 billion, as 117 IPOs doubled in price on their first day of trading.¹⁶ Almost all of these cases involved very young firms, many of which had venture capital backing. The 466 IPOs meeting our sample criteria raised \$65 billion in aggregate proceeds and paid \$4 billion in gross spreads. By contrast, in the entire 1975–1994 period, only 10 IPOs with an offer price of more than \$5.00 per share doubled on their first day. During 1995–1998, 29 IPOs doubled. Thus the patterns that we seek to explain have not been going away. Indeed, the topic of why issuers aren't upset about the amount of money left on the table is more topical than ever. Every month from November 1998 to March 2000 saw average first-day returns in excess of 30%, leaving no doubt that hot issue markets continue to exist. The year 2000 saw another \$27 billion left on the table, as 77 more IPOs doubled in price on the first day.

¹⁶ Jay Ritter's Web site, bear.cba.ufl.edu/ritter, includes a current listing of IPOs that have doubled in price on the offer day, as well as the largest IPOs ranked by the amount of money left on the table.

Appendix: Autocorrelation of Monthly Average First-Day Returns

Following the arguments made in Section 6, let's start with Equations (1) and (2):

$$OP_i = MID_i + \alpha_1 \cdot r_{0,T_i}^{\text{market}} \cdot MID_i + \alpha_2 \cdot r_{0,T_i}^{\text{public}} \cdot MID_i + \alpha_3 \cdot r_{0,T_i}^{\text{private}} \cdot MID_i + \tau_i \quad (1)$$

$$P_i = 1.14 \cdot [MID_i + \beta \cdot r_{0,T_i}^{\text{market}} \cdot MID_i + \gamma_2 \cdot r_{0,T_i}^{\text{public}} \cdot MID_i + \gamma_3 \cdot r_{0,T_i}^{\text{private}} \cdot MID_i] + \varepsilon_i \quad (2)$$

All definitions of variables in these two equations are the same as in the article, with MID being the midpoint of the original file price range. $r_{0,T_i}^{\text{public}}$ could be interpreted as industry movements above and beyond market movements. Assume that $r_{0,T_i}^{\text{public}}$ and $r_{0,T_i}^{\text{market}}$ are independent, as are $r_{0,T_i}^{\text{private}}$ and $r_{0,T_i}^{\text{market}}$. We can then simplify Equations (1) and (2) as

$$OP_i = MID_i + \alpha_1 \cdot r_{0,T_i}^{\text{market}} \cdot MID_i + \tau_i' \quad (1')$$

$$P_i = 1.14 \cdot [MID_i + \beta \cdot r_{0,T_i}^{\text{market}} \cdot MID_i] + \varepsilon_i' \quad (2')$$

The error terms capture not only firm-specific errors, but also valuation-related information that is both private and public (such as recent industry movements). Consequently the error terms in Equations (1') and (2') will be correlated. Divide both sides of the above equations by MID_i to get

$$\frac{OP_i}{MID_i} = 1 + \alpha_1 \cdot r_{0,T_i}^{\text{market}} + \tau_i'' \quad (1'')$$

$$\frac{P_i}{MID_i} = 1.14 + 1.14\beta \cdot r_{0,T_i}^{\text{market}} + \varepsilon_i'' \quad (2'')$$

where $\tau_i'' = \tau_i'/MID_i$ and $\varepsilon_i'' = \varepsilon_i'/MID_i$. Then the first day return for IPO i is

$$R_i = \frac{P_i - OP_i}{OP_i} = \frac{.14 + (1.14\beta - \alpha_1) \cdot r_{0,T_i}^{\text{market}} + \varepsilon_i'' - \tau_i''}{1 + \alpha_1 \cdot r_{0,T_i}^{\text{market}} + \tau_i''} \quad (3)$$

To simplify the problem and make it tractable, we assume that a linear model is a good approximation. That is, we assume Equation (3') is good enough to capture the relationship between R_i and $r_{0,T_i}^{\text{market}}$ specified in Equation (3):

$$R_i = \lambda_0 + \lambda_1 \cdot r_{0,T_i}^{\text{market}} + \eta_i \quad (3')$$

Since the monthly average of IPO first-day returns is a linear combination of all IPO first-day returns, let's first derive the formula for the covariance of first-day returns of two IPOs, namely i and j . Using Equation (3'), we know

$$\text{cov}(R_i, R_j) = \text{cov}(\lambda_0 + \lambda_1 \cdot r_{0,T_i}^{\text{market}} + \eta_i, \lambda_0 + \lambda_1 \cdot r_{0,T_j}^{\text{market}} + \eta_j) \quad (4)$$

Assume that, except for common market movements, the pricing processes are independent between any two IPOs. (This assumption ignores industry factors, and thus understates the covariance.) Hence the error terms, η_i for all i and j , are also independent. Consequently Equation (4) becomes

$$\text{cov}(R_i, R_j) = \lambda_1^2 \text{cov}(r_{0,T_i}^{\text{market}}, r_{0,T_j}^{\text{market}}) \quad (4')$$

Now let's consider the monthly average IPO first-day return. To make the case as simple as possible, assume that there are 20 trading days per month, and that there is only one IPO each trading day. Denote the monthly average as \bar{R}_k for month k . Then

$$\bar{R}_k = \frac{1}{20} \sum_{i=1}^{20} R_i \quad (5)$$

where i denotes the i th trading day of month k .

For two consecutive months, month k and month $k + 1$, the covariance between the IPO first-day return averages of these two months is then as follows

$$\begin{aligned} \text{cov}(\bar{R}_k, \bar{R}_{k+1}) &= \text{cov}\left(\frac{1}{20} \sum_{i=1}^{20} R_i, \frac{1}{20} \sum_{j=1}^{20} R_j\right) \\ &= \frac{1}{400} \sum_{i=1}^{20} \sum_{j=1}^{20} \text{cov}(R_i, R_j) \\ &= \frac{\lambda_1^2}{400} \sum_{i=1}^{20} \sum_{j=1}^{20} \text{cov}(r_{0,T_i}^{\text{market}}, r_{0,T_j}^{\text{market}}). \end{aligned} \tag{6}$$

Please notice that j denotes the j th trading day in month $k + 1$.

Let's look more closely at $\text{cov}(r_{0,T_i}^{\text{market}}, r_{0,T_j}^{\text{market}})$. Notice that those two market factors in the covariance term are nothing more than the cumulative market returns during the 15-day window (from day -14 to day 0 when the IPO went public) of IPO i and IPO j , respectively. That is,

$$r_{0,T_i}^{\text{market}} = \sum_{m=0}^{14} r_{T_i-m} \tag{7}$$

$$r_{0,T_j}^{\text{market}} = \sum_{n=0}^{14} r_{T_j-n}, \tag{8}$$

where r_{T_i-m} and r_{T_j-n} are simply daily continuously compounded market returns on day $T_i - m$ and $T_j - n$, respectively. T_i denotes the i th trading day when IPO i went public and T_j is the j th trading day when IPO j went public. Given Equations (7) and (8), we then have

$$\text{cov}(r_{0,T_i}^{\text{market}}, r_{0,T_j}^{\text{market}}) = \text{cov}\left(\sum_{m=0}^{14} r_{T_i-m}, \sum_{n=0}^{14} r_{T_j-n}\right). \tag{9}$$

To simplify the calculation, let's assume $\text{cov}(r_i, r_j) = 0$ for any $i \neq j$, and $\text{var}(r_i) = \sigma^2$ for all i . Then Equation (9) becomes

$$\text{cov}(r_{0,T_i}^{\text{market}}, r_{0,T_j}^{\text{market}}) = D_{ij} \cdot \sigma^2, \tag{10}$$

where D_{ij} is the number of days r_{0,T_i} and r_{0,T_j} have in common.

Given Equation (10), and that i and j denote the i th and j th trading days in month k and month $k + 1$, respectively, we can easily get

$$\begin{aligned} \text{cov}(\bar{R}_k, \bar{R}_{k+1}) &= \frac{\lambda_1^2}{400} \sum_{i=1}^{20} \sum_{j=1}^{20} \text{cov}(r_{0,T_i}^{\text{market}}, r_{0,T_j}^{\text{market}}) \\ &= \frac{\lambda_1^2 \sigma^2}{400} \sum_{i=1}^{20} \sum_{j=1}^{20} D_{ij}. \end{aligned} \tag{11}$$

Since what we care about is the autocorrelation of the IPO first-day return monthly average, Equation (11) only gives us the numerator part. Now let's turn to the denominator part, the variance of the monthly average IPO first-day returns.

Assume $\text{var}(\bar{R}_k) = \text{var}(\bar{R}_{k+1})$. Notice that

$$\begin{aligned}
 \text{var}(\bar{R}_k) &= \text{var}\left(\frac{1}{20} \sum_{m=1}^{20} R_m\right) \\
 &= \text{cov}\left(\frac{1}{20} \sum_{m=1}^{20} R_m, \frac{1}{20} \sum_{n=1}^{20} R_n\right) \\
 &= \frac{1}{400} \sum_{m=1}^{20} \sum_{n=1}^{20} \text{cov}(R_m, R_n) \\
 &= \frac{\lambda_1^2}{400} \sum_{m=1}^{20} \sum_{\substack{n=1 \\ m \neq n}}^{20} \text{cov}(r_{0,T_m}^{\text{market}}, r_{0,T_n}^{\text{market}}) + \frac{1}{400} \sum_{n=1}^{20} \text{var}(R_n). \tag{12}
 \end{aligned}$$

Please note that, different from Equation (11), m and n in Equation (12) denote the m th and n th trading days in the same month, month k . Similar to the covariance term between two consecutive months, the first term in Equation (12) can also be expressed in terms of the variance of the market return. That is,

$$\text{var}(\bar{R}_k) = \frac{\lambda_1^2 \sigma^2}{400} \sum_{m=1}^{20} \sum_{\substack{n=1 \\ m \neq n}}^{20} D_{mn} + \frac{1}{400} \sum_{n=1}^{20} \text{var}(R_n). \tag{13}$$

Now let's look more closely at the second term in Equation (13). From Equation (3'), we have

$$\text{var}(R_n) = \text{var}(\lambda_0 + \lambda_1 \cdot r_{0,T_n}^{\text{market}} + \eta_n) = 15\lambda_1^2 \sigma^2 + \sigma_\eta^2, \tag{14}$$

where σ_η^2 is the variance of the error term, η_n . Given Equation (14), Equation (13) becomes

$$\text{var}(\bar{R}_k) = \frac{\lambda_1^2 \sigma^2}{400} \sum_{m=1}^{20} \sum_{\substack{n=1 \\ m \neq n}}^{20} D_{mn} + \frac{300}{400} \lambda_1^2 \sigma^2 + \frac{1}{20} \sigma_\eta^2. \tag{15}$$

Given Equations (12) and (15), the first-order autocorrelation of monthly average IPO first-day returns, denoted as ρ , is as follows

$$\begin{aligned}
 \rho &= \frac{\frac{\lambda_1^2 \sigma^2}{400} \sum_{i=1}^{20} \sum_{j=1}^{20} D_{ij}}{\frac{\lambda_1^2 \sigma^2}{400} \sum_{m=1}^{20} \sum_{\substack{n=1 \\ m \neq n}}^{20} D_{mn} + \frac{300}{400} \lambda_1^2 \sigma^2 + \frac{1}{20} \sigma_\eta^2} \\
 &= \frac{\sum_{i=1}^{20} \sum_{j=1}^{20} D_{ij}}{\sum_{m=1}^{20} \sum_{n=1}^{20} D_{mn} + \frac{20\sigma_\eta^2}{\lambda_1^2 \sigma^2}} \\
 &= \frac{560}{3380 + \frac{20\sigma_\eta^2}{\lambda_1^2 \sigma^2}}. \tag{16}
 \end{aligned}$$

Now let's plug in some numbers. Let $\sigma^2 = 0.0001$, $\lambda_1 = 1.33$ (as in Table 5), and $\sigma_\eta^2 = 0.01$. Then $\rho = 0.124$.

Sensitivities to this base case are easily performed. If there are 40 IPOs per month rather than 20, Equation (16) yields $\rho = 0.142$. This increase is because the firm-specific volatility is diversified away in the monthly portfolio return. If the standard deviation of the first-day returns increases from 10% per IPO to the 23.8% in our Table 2, then $\rho = 0.057$. This decrease in the autocorrelation is because of the greater noise in the monthly average. If the length of the

lagged return interval increases from 15 trading days to 21 trading days (the day of the offering plus the prior four weeks), then $\rho = 0.224$.

These results demonstrate that as long as the offer price does not fully adjust to publicly available information about prior market returns, the monthly average first-day returns will be autocorrelated. If there are important industry factors to which there is partial adjustment, the autocorrelation will be even stronger. The empirical autocorrelation coefficient that we have estimated is 0.50, so either industry effects are important, or there is some other source of autocorrelation that we have not modeled (or both).

References

- Aggarwal, R., 2000, "Stabilization Activities by Underwriters After Initial Public Offerings," *Journal of Finance*, 55, 1075–1103.
- Atchison, M. D., K. C. Butler, and R. R. Simonds, 1987, "Nonsynchronous Security Trading and Market Index Autocorrelation," *Journal of Finance*, 42, 111–118.
- Baron, D., 1982, "A Model of the Demand of Investment Banking Advising and Distribution Services for New Issues," *Journal of Finance*, 37, 955–976.
- Barry, C., 1989, "Initial Public Offerings Underpricing: The Issuer's View—A Comment," *Journal of Finance*, 44, 1099–1103.
- Beatty, R. P., and J. R. Ritter, 1986, "Investment Banking, Reputation, and the Underpricing of Initial Public Offerings," *Journal of Financial Economics*, 15, 213–232.
- Benveniste, L., and P. Spindt, 1989, "How Investment Bankers Determine the Offer Price and Allocation of New Issues," *Journal of Financial Economics*, 24, 343–361.
- Benveniste, L., A. Ljungqvist, W. Wilhelm, and X. Yu, 2001, "Evidence of Information Spillovers in the Production of Investment Banking Services," working paper, University of Minnesota.
- Brealey, R., and S. Myers, 1996, *Principles of Corporate Finance* (5th ed.), McGraw-Hill, New York.
- Chan, L., and J. Lakonishok, 1992, "Robust Measurement of Beta Risk," *Journal of Financial and Quantitative Analysis*, 27, 265–282.
- Chen, H., and J. R. Ritter, 2000, "The Seven Percent Solution," *Journal of Finance*, 55, 1105–1131.
- Clarkson, P. M., and R. Thompson, 1990, "Empirical Estimates of Beta When Investors Face Estimation Risk," *Journal of Finance*, 45, 431–453.
- Derrien, F., and K. L. Womack, 2002, "Auctions vs. Book-Building and the Control of Underpricing in Hot IPO Markets," forthcoming in *Review of Financial Studies*.
- Dunbar, C., 2000, "Factors Affecting Investment Bank Initial Public Offering Market Share," *Journal of Financial Economics*, 55, 3–41.
- Fitzgibbon, J. L., Jr., 1998, "The Last Six Months," *IPO Reporter*, July 6, 22.
- Habib, M., and A. Ljungqvist, 2001, "Underpricing and Entrepreneurial Wealth Losses: Theory and Evidence," *Review of Financial Studies*, 14, 433–458.
- Hanley, K. W., 1993, "Underpricing of Initial Public Offerings and the Partial Adjustment Phenomenon," *Journal of Financial Economics*, 34, 231–250.
- Hanley, K. W., A. Kumar, and P. Seguin, 1993, "Price Stabilization in the Market for New Issues," *Journal of Financial Economics*, 34, 177–197.
- Hof, R. D., 1999, "Inside an Internet IPO," *Business Week*, Sept. 6, 60–73.
- Ibbotson, R. G., and J. F. Jaffe, 1975, "'Hot issue' markets," *Journal of Finance*, 30, 1027–1042.

- Ibbotson, R. G., J. Sindelar, and J. R. Ritter, 1994, "The Market's Problems with the Pricing of Initial Public Offerings," *Journal of Applied Corporate Finance*, 7, 66–74.
- Jenkinson, T., and A. Ljungqvist, 2001, *Going Public: The Theory and Evidence on How Companies Raise Equity Finance*, Clarendon Press, Oxford.
- Kahneman, D., and A. Tversky, 1979, "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 263–291.
- Krigman, L., W. Shaw, and K. Womack, 1999, "The Persistence of IPO Mispricing and the Predictive Power of Flipping," *Journal of Finance*, 54, 1015–1044.
- Krigman, L., W. Shaw, and K. Womack, 2001, "Why Do Firms Switch Underwriters?," *Journal of Financial Economics*, 60, 245–284.
- Logue, D., 1973, "On the Pricing of Unseasoned Equity Issues: 1965–69," *Journal of Financial and Quantitative Analysis*, 8, 91–103.
- Logue, D., R. Rogalski, J. Seward, and L. Foster-Johnson, 2001, "What's Special About the Role of Underwriter Reputation and Market Activities in IPOs?," forthcoming in *Journal of Business*.
- Lowry, M., and G. W. Schwert, 2002, "IPO Market Cycles: Bubbles or Sequential Learning?," forthcoming in *Journal of Finance*.
- Lucchetti, A., 1999, "SEC Probes Rates Funds Firms Pay for Commissions," *Wall Street Journal*, Sept. 16, C1.
- Maksimovic, V., and H. Unal, 1993, "Issue Size Choice and 'Underpricing' in Thrift Mutual-to-Stock Conversions," *Journal of Finance*, 48, 1659–1692.
- Mandelker, G., and A. Raviv, 1977, "Investment Banking: An Economic Analysis of Optimal Underwriting Contracts," *Journal of Finance*, 32, 683–694.
- Ritter, J., 1991, "The Long-Run Performance of Initial Public Offerings," *Journal of Finance*, 46, 3–27.
- Shefrin, H. M., and M. Statman, 1984, "Explaining Investor Preference for Cash Dividends," *Journal of Financial Economics*, 13, 253–282.
- Sternberg, T., 1989, "Bilateral Monopoly and the Dynamic Properties of Initial Public Offerings," working paper, Vanderbilt University.
- Thaler, R. H., 1980, "Towards a Positive Theory of Consumer Choice," *Journal of Economic Behavior and Organization*, 1, 39–60; reprinted in Thaler, R. H., 1991, *Quasi Rational Economics*, Russell Sage Foundation, New York.
- Uttal, B., 1986, "Inside the Deal that Made Bill Gates \$350,000,000," *Fortune*, 24, 343–361.
- Welch, I., 1992, "Sequential Sales, Learning, and Cascades," *Journal of Finance*, 47, 695–732.
- White, H., 1980, "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," *Econometrica*, 48, 817–838.

