

Stocks as Money: Convenience Yield and the Tech-Stock Bubble.

John H. Cochrane¹

May 23, 2002

¹Graduate School of Business, University of Chicago, 1101 E. 58th St. Chicago IL 60637. This paper was prepared for the Federal Reserve Bank of Chicago Conference on Asset Price Bubbles, April 22-24, 2002. I thank Thomas Chevrier for research assistance, Owen Lamont for data, extensive comments, and many helpful discussions, and Matthew Richardson for data. Revised versions of this paper (and this version with color graphs) can be found at <http://gsbwww.uchicago.edu/fac/john.cochrane/research/Papers/>.

Abstract

What caused the rise and fall of tech stocks? I argue that a mechanism much like the transactions demand for money drove many stock prices above the “fundamental value” they would have had in a frictionless market. I start with the Palm/3Com microcosm and then look at tech stocks in general. High prices are associated with high volume, high volatility, low supply of shares, wide dispersion of opinion, and restrictions on long-term short selling. I review competing theories, and only the convenience yield view makes all these connections.

1 Introduction

What caused the rise and fall of tech stocks in the late 1990s? I suggest that a mechanism much like the transactions demand for money drove many stock prices above the “fundamental value” they would have had in a frictionless market.

During the boom, there was an intense demand for short-term trading in tech stocks. As a result of market frictions, such trading requires shares of the stock – if no shares are outstanding, there’s no way to bet one way or the other on the future of a company. Few shares were available for trading, so the available shares gave a convenience yield: People were willing to hold them for a little while for short-term trading, even though they knew that the shares were overvalued as a long-term investment, just as people will briefly hold money even though it depreciates rapidly in a hyperinflation.

As Ofek and Richardson (2001) document, tech stocks fell when many more shares became available, due to a combination of IPOs, expiration of lockup periods, and increasing ability to sell short, while at the same time the speculative demand for shares mirrored in share volume declined dramatically. As increasing money supply and declining transactions demand lead to lower interest rates – money less overpriced relative to bonds – these events sharply reduced the convenience yield of shares.

This paper simply documents the analogy between tech stocks and conventional money demand. I start with a microcosm, the 3Com/Palm event, and then I extend the lessons of that microcosm to the Nasdaq / tech stock experience as a whole. I verify that the elements of a trading-related convenience yield are there in each case, in particular that high prices are associated with high volume and low share supply. I conclude with a review of various theories. The key point is that in the tech stock boom and bust, as in the famous historical “bubbles,” high prices come along with a trading frenzy. *None* of the alternative theories says anything about this linkage – they can predict high prices just as easily with no volume. The convenience yield inextricably links the price rise and decline with the rise and decline of trading.

This paper is an interpretive review. Most of the empirical work is either taken directly from or closely inspired by the work of Lamont and Thaler (2001) and Ofek and Richardson (2001, 2002). My interpretation of the evidence is quite different.

2 3Com, Palm and Convenience Yield

3Com and Palm

On March 2, 2000, 3Com sold 5% of its shares of Palm in an initial public offering. It retained 95% of the shares, and announced that it would give those shares to 3Com shareholders by the end of the year. Each 3Com share would get approximately 1.5 shares of Palm. (Most of the data and facts about this event come from Lamont and

Thaler 2001.)

There were two ways to end up with a share of Palm at the end of 2000: you could buy one share of Palm directly, or you could buy 1/1.5 shares of 3Com. At the end of trading on March 2, a share of Palm bought directly cost \$95.06. 3Com closed at \$81.81, so a share of Palm bought by buying 3Com cost $\$81.81/1.5 = \54.54 – a much lower price for an apparently identical security, and you get the rest of 3Com for free.

Figure 1 plots daily data on the price of Palm stock and 1/1.5 times the price of 3Com stock. As Figure 1 shows, it was cheaper to buy Palm “implicitly” by buying 3Com than it was to buy it directly through mid-May. The prices in Figure 1 imply that the *rest* of 3Com (the “stub”) was valued by the market at a *negative* amount – minus 22 billion dollars at the end of the day on March 2. (The sharp drop in 3Com in late July seen in Figure 1 comes on the day that it spun off its remaining Palm shares. The market apparently had no problems adding and subtracting on that day!)

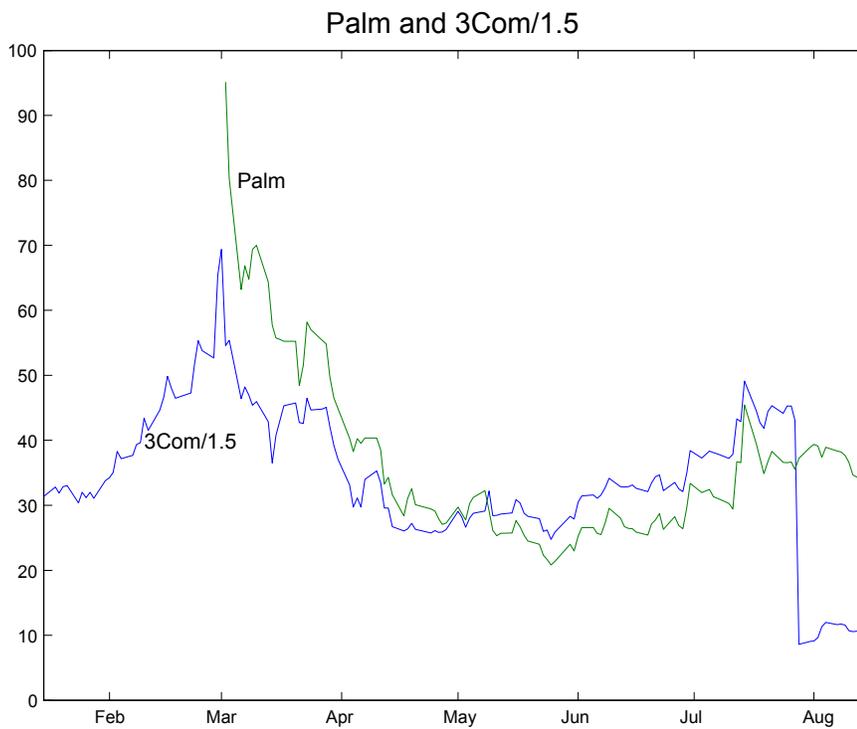


Figure 1: Price of Palm and price of 3Com/1.5 from Palm’s IPO to the eventual spinoff.

This event seems a clear violation of the law of one price. 3Com should always be worth *at least as much* as its holdings of Palm.

This event is an interesting microcosm in which to start thinking about the stock

market events of the end of the decade. The value of Palm embedded in 3Com is an easily-measured lower bound on the “fundamental value,” so this event allows us pretty cleanly to look at a case of a security (Palm) whose price was above such a “fundamental value.” Then, we can see to what extent the same lessons might apply to other stocks, and tech or the NASDAQ index, for which “fundamental value” measures are much harder to estimate.

Similar events, obvious objections.

The 3Com/Palm event was not isolated. Lamont and Thaler document 6 additional carve-outs with negative stub values in the 1996-2000 period. Mitchell, Pulvino and Stafford (2002) find 82 cases in a longer sample in which the implied value of a parent company is less than the value of its holdings of a publicly traded subsidiary. More generally, there have been many puzzling circumstances in which a rapid rise in the stock price of a partially owned subsidiary does not affect the parent’s stock price. For example, in 1999 GM had issued tracking stock for its Hughes Electronics unit, and also had a 20% stake in publicly traded Commerce One. (The facts are from Lamont 2000). Between September 1999 and January 2000, Hughes stock rose 97 percent and Commerce One stock rose 413 percent. GM’s stock was barely affected. Lamont cites analyst calculations that this move left GM’s auto business a price/earnings ratio of only 1.5, at the same time Ford’s price/earnings ratio was 7 and DaimlerChrysler’s was 12. The value of the rest of GM did not fall below zero, but the frictionless model is on thin ice if we have to assume that shocks to GM’s fundamentals are strongly *negatively* correlated with shocks to Hughes and Commerce One fundamentals, and uncorrelated with those of Ford and DaimlerChrysler.

These and related observations suggest to many observers that a downward sloping “demand for shares” is at work. The prices of the *available* shares are being set as if the unavailable shares – Palm Shares held by 3Com, Hughes shares held by GM, tech stock shares held by insiders, etc. – did not exist.

Why aren’t such price differences arbitrated away? The Lamont and Thaler, Mitchell Pulvino and Stafford and other investigations of these events carefully document the institutional details that prevent arbitrage. In the real world, you cannot costlessly short Palm and buy 3Com shares, in anticipation of your arbitrage profits after spinoff. Short sales require you to borrow stock before you sell it. In the 3Com/Palm case, 3Com stock was often simply not available for borrowing. If your broker could find some, you may have had to pay dearly for the privilege, unlike the textbook case that you receive interest on the proceeds of the short sale. Furthermore, if the spread widens, you may be wiped out before the price finally rights itself. The short loan may be called, and you may be unable to reestablish the short position – short loans must be reestablished *daily*.

In the end though, the fact that it could not be arbitrated away does not resolve the puzzle. The puzzle is, *Why are Palm and 3Com prices different in the first place? Who is buying overpriced Palm shares and why? What is the source of the “demand for shares?”*

All results in finance are controversial, and this one is as well. One can quarrel whether it really was an arbitrage. True, fundamental, stub values can be negative. Though stock is a limited liability security, stubs are not. $P \geq 0$ and $P + C \geq 0$ do not imply $C \geq 0$. For example, 3Com may borrow, using Palm shares as collateral, and then go bankrupt. In addition, the spinoff is not 100% sure to happen. 3Com can postpone it or cancel it (as they can and did advance it from December to July). 3Com may be acquired, and the new parent may cancel the spinoff. Note in Figure 1 the sudden end of the negative stub value, on the day that the IRS approved the tax-free status of the spinoff. This fact alone suggests that there is at least some real risk that the spinoff not happen. Mitchell, Pulvino and Stafford discuss all the ways in which apparent negative stub arbitrages can fail to work out, and find that a surprising 30% of their negative stub “arbitrage opportunities” terminate without removing the mispricing. This story is strained for 3Com/Palm. It’s hard to imagine that the correct valuation of 3Com less Palm was negative 22 billion dollars in March, and then recovered steeply to the spinoff in July.

T bills and dollar bills

On March 2, 2000, the one-year treasury bill sold for¹ \$94.17. A treasury bill is, of course, a sure promise to receive \$100 in a year’s time. One could also get a claim to \$100 in a year’s time by buying \$100 directly and holding it for a year. *Two ways of getting exactly the same payoff have a different price.* Why does anybody hold an overpriced dollar (Palm) when they could hold a cheaper Treasury bill (3Com) instead?

We have lots of good stories for this clearcut violation of the law of one price makes sense. We don’t need an irrational, psychological, or behavioral attachment to “dollars” rather than “treasury bills” to explain it. People hold money because it provides “liquidity services” that make up for its poor rate of return. A poor rate of return is the same thing as a “too high” price. Much of the point of monetary economics is devoted to explaining “rate of return dominance.”

This analogy is not just creative residual-naming. Monetary economics makes some quite sharp predictions about when this “mispricing” can occur. If Palm/3Com is “like” dollars/treasury bills, we must see the standard predictions of money demand hold true.

2.1 Money-like facts about 3Com and Palm

1. Huge turnover and small short term losses

Nobody holds dollars today as a way of getting dollars a year from now. (Except drug dealers and so on who value the anonymity of dollars, but the traditional theory of money demand is the right analogy for stocks). We each hold dollars for a short

¹Data from the Federal Reserve H.15 release, <http://www.federalreserve.gov/releases/h15/data/b/tbsm1y.txt>

time, between trips to the bank and transactions. We're happy to hold dollars for these short times, despite the fact that they depreciate relative to Treasury bills, because we need to hold dollars briefly. Dollars turn over quickly, and as interest rates rise – as dollars rise in price relative to treasury bills – dollars turn over more quickly.

Palm turned over quickly. Lamont and Thaler (Table 8) report that on average 19% of the available Palm shares changed hands *every day* in the 20 days after the IPO. The even more dramatic case of Creative/Ubic had a 106% average daily turnover, and Lamont and Thaler's 6 cases average 38% *daily* turnover (Table 8). Figure 2 plots Palm daily turnover. As you can see, there are many days with huge turnover. Volume on the first day was 1.5 times the total issue, as prices fell quickly from as much as \$160 to the \$92 close. By comparison, only 4.5% of 3Com shares were traded every day in the same period, and daily share volumes for typical stocks are 2% or less. This means that on average, Palm shareholders held the stock for less than 5 days², and during peak periods average holdings were much shorter than that!

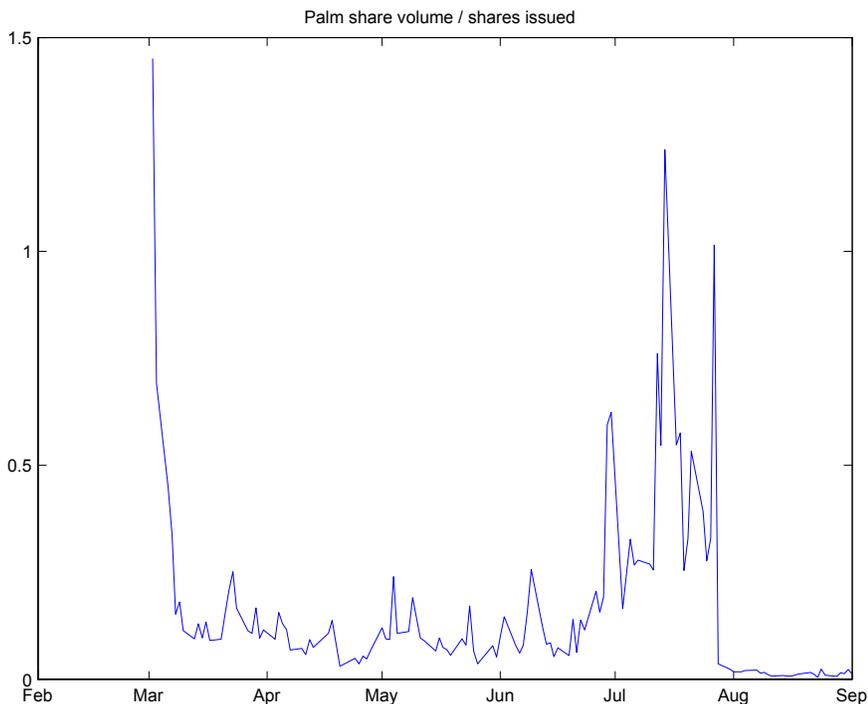


Figure 2: Palm daily turnover (volume / shares issued). The sharp fall in late July comes when the remaining 95% of shares are spun off by 3Com.

Lamont and Thaler conclude that people who bought Palm rather than 3Com were “irrational” and “just making a mistake.” They could have gotten Palm a lot

²NASDAQ volume includes dealer trades, so one might argue that the correct number is 10 days rather than 5. On the other hand dealers are people too. More generally, there was surely great heterogeneity in holding times.

cheaper by buying 3Com and waiting for the spinoff; a massive version of buying the name brand rather than the generic. But these turnover rates suggest that few Palm buyers did, in fact, make that mistake, and the vast majority of Palm shares were held for very short times.

At a 5 day horizon, holding Palm rather than 3Com is not so obviously stupid. Take at face value that the Palm share at \$95.06 will decline to the implicit \$54.54 value from buying 3Com in 9 months. This means a -42.6% relative return over 9 months – pretty bad. But it is only a negative $2/10$ percent *daily* return, or a one percent negative *5 day* return. Now, losing on average $2/10$ of a percent as a day trader, or one percent as a 5-day trader isn't *bright*, but it's much further from idiotic than the huge loss of buying Palm rather than 3Com and holding it for a year.

Palm stock was also tremendously volatile during this period, with 7.15% standard deviation of daily returns and 15.4% standard deviation of 5 day returns. The latter is about the same as the volatility of the S&P500 index over an entire *year*. Figure 3 plots the distribution of 1 and 5 day Palm returns. Imagine that you are betting on the movements of Palm over a few days. Clearly the fact that Palm will on average drift down two tenths of a percent per day is completely drowned out by the typical movements. Only a small bit of information about the short-term movements will swamp the information of a $-2/10\%$ daily drift down, and make Palm a smart investment. The drift is a small loss, of the same order of magnitude as the bid/ask spread, commissions, or the loss from taking a short term rather than long term capital gain. The Good Housekeeping guide to careful investing warns you to minimize these losses, but we usually don't jump to "irrationality" to explain why people are a bit sloppy about managing small losses, especially when quick decisions can bring such large gains.

2. Price and volume

A crucial prediction of monetary economics is that velocity increases at higher interest rates, or as the spread between the value of a dollar and a T-bill increase. As interest rates rise, people pay more attention to economizing cash holdings, and so turnover increases. If some sort of transactions demand is behind "overvalued" Palm shares, then the "overvaluation" should be associated with huge volume.

We already know that Palm's volume was in fact huge on average. Figure 4 shows that time variation in Palm volume also lines up with time-variation in Palm's prices. (The plot presents a 5 day centered moving average of share volume. Volume varies enormously from day to day. Prices incorporate expected returns for a long time in the future; so prices will be associated with longer-run movements in volume. Dollar volume is the economically more meaningful measure, but I plot share volume to emphasize that the positive correlation between volume and price does not just come from a constant share volume multiplied by varying prices.)

A positive correlation between price and volume is not a common stylized fact of the market microstructure literature. For example, Gallant, Rossi and Tauchen's

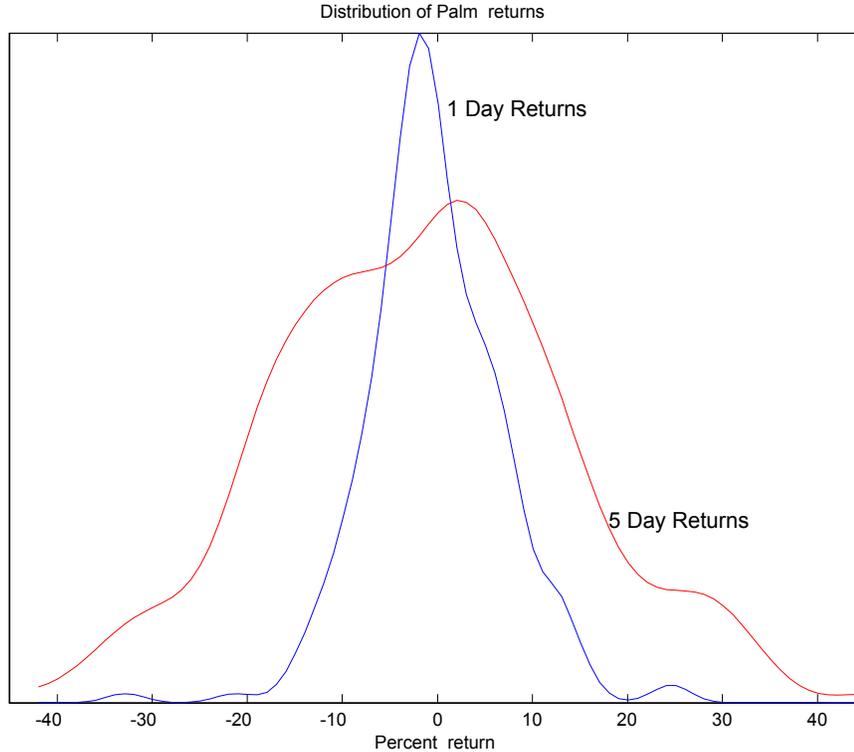


Figure 3: Distribution of Palm daily returns, 2002. The smoothed histogram uses a normal distribution window with a 2% standard deviation window width.

(1992) comprehensive study lists four stylized facts, starting with the correlation of volume with volatility, but do not mention a broad-based correlation of volume with price. It isn't unknown – for example, Brennan and Subrahmanyam (1996) and Jones (2001, Table 2) find that high share turnover (and other liquidity measures) forecast low subsequent returns, which is the same thing as a high price – but it isn't one of the most commonly studied effects. The spike in volume surrounding the sharp price decline in late December is a more typical volume event. The spike in volume around mid-July comes contemporaneous to the 3Com spinoff, and is therefore also unlikely to be primarily driven by a transactions demand. It is significant that a positive correlation between price and volume does show up. It suggests an unusual case in which the transactions-based relation between price and volume can stick out above all the other more usual effects.

3. Arbitrage and short sales constraints

As we have seen, a crucial feature of Palm/3Com was that short selling was at first impossible, then very expensive.

A restriction on short selling is vital to maintaining the dollar/treasury bills spread as well. Why doesn't arbitrage remove the price difference between dollars and treasury bills? The arbitrage is to short dollar bills and buy treasury bills. Alas, you can't

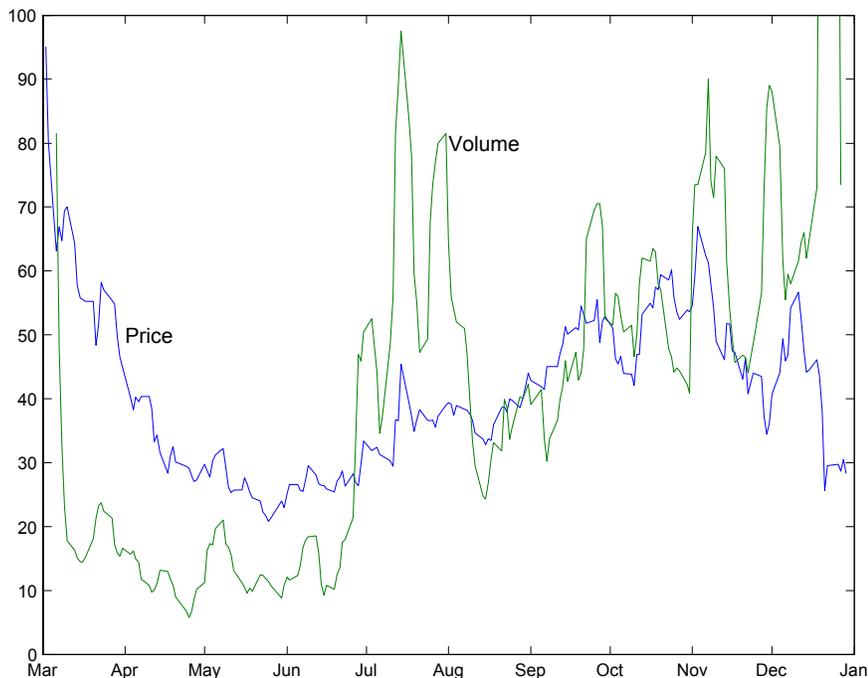


Figure 4: Palm price and a 5 day centered moving average of share volume.

short dollar bills—printing them is illegal. Printing close substitutes – banknotes, or small denomination bearer bonds – is also illegal.

4. *Share supply and short sales*

Money is more overpriced – interest rates are higher – when money supply is lower. In this context it is interesting that the Palm “overvaluation” happened while only 5% of Palm stock was outstanding, and 95% retained by 3Com. If you wanted to bet on Palm computers, up or down, you had to compete for one of the very small number of shares outstanding. (Actually, what matters is the “float,” the number of shares easily available for trading. Many shares are not actively traded or available to be lent for short even though in private hands. Low float can come from a small amount issued, or from a small amount of a large issue available for trading. Unfortunately, we don’t have a clean definition and even less data on float.)

Short selling can act like inside money (bank account) creation as a way to increase the supply of shares. If A lends shares to B and B sells to C, then both A and C have long positions even though there is only one share outstanding.

Despite the costs of short positions, Palm short sales were massive, increasing steadily to 147% of available shares in July. That means that on average each share was bought, lent to short, sold and then half were bought, lent, and shorted again. (After the spinoff, the total number of Palm shares sold short did not change, but the supply of available shares jumped so that the fractions returned to normal. Inter-

estingly, Lamont and Thaler show that *parents* averaged 3.7% short interest. While much lower than the subsidiaries, it is a testament to the divergence of opinion on tech stocks and the emphasis on high-frequency trading that anyone was short the long end of an arbitrage opportunity!)

Figure 5 plots Palm’s “overpricing” relative to 3Com, i.e. $\text{Palm}/(3\text{Com}/1.5) - 1$, against the share supply induced by shorting, $1 + \text{short interest}/\text{shares issued}$. You can see the nice pattern – as share supply increased, the “overpricing” decreased steadily.

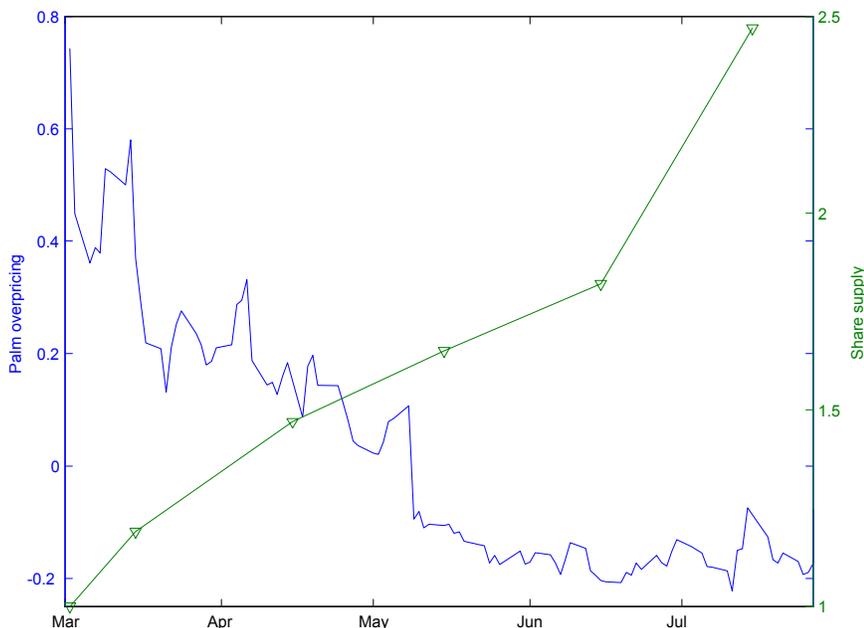


Figure 5: Palm overpricing and share supply. “Palm overpricing” is $\text{Palm price}/(3\text{Com price}/1.5) - 1$. “Share supply” is $1 + \text{short interest}/\text{shares issued}$.

It is a little surprising that Palm prices did not fall when the spinoff occurred and a huge new volume of shares was available for trading. However, this event was widely anticipated, so the price decline would have to be slow and hard to measure. The explosion in shorting in July indicates many people expected such a price decline, and the massive short interest may have gone a long way to providing the required extra share supply before the spinoff.

5. *Poor Substitutes – if you want to trade, you have to trade Palm*

There can only be a money premium if money is a “special asset,” if there are no good substitutes such as private banknotes. There can only be a convenience yield for Palm stock if there are no good substitutes for the purposes of high-frequency trading. Are there?

3Com and Palm are poorly linked. The obvious substitute for Palm is 3Com, since 3Com holds 95% of Palm stock. If you want to bet on Palm for a few days, why not buy 3Com instead and save the 1% negative drift? Surely, any news that impacts Palm’s prospects will affect 3Com as well.

Table 1 evaluates this option. The first row runs Palm returns on 3Com returns. This is simple, but the two returns may have a spuriously high correlation as they will both rise when the market rises. The right strategy for betting on Palm is to hedge the Palm investment with an investment in the market index, for a beta zero strategy that bets only on Palm’s fortunes. Thus, in the second row I estimate a beta on the NASDAQ index for each of Palm and 3Com, and then run a regression of the Palm market model residual on the 3Com market model residual. The bottom set of rows evaluates the Palm / 3Com return correlation at a 5 day horizon. We know Palm and 3Com will converge at a 6 month horizon, perhaps the delinking at a one day horizon disappears by the typical 5 day horizon.

Table 1 shows that buying or selling 3Com is a very poor way of betting on Palm at short horizons. While the regression coefficients b are gratifyingly close to one, the R^2 are surprisingly low, given that *all* of 3Com’s value (and then some!) is attributable to its Palm shares. Of the 7% daily Palm volatility and 13-15% five day Palm volatility, 4.5% daily Palm volatility and 8-10% 5 day Palm volatility are missed by the strategy of buying 3Com in place of Palm.

	b	R^2	$\sigma(y)$	$\sigma(\varepsilon)$
One day returns:				
$\text{Palm}_t = a + b3\text{Com}_t + \varepsilon_t$	0.96	0.60	7.2	4.5
$(\text{Palm}_t - \beta\text{Nasdaq}_t) = a + b(3\text{Com}_t - \beta\text{Nasdaq}_t) + \varepsilon_t$	0.93	0.53	6.9	4.6
Five day returns:				
$\text{Palm}_t = a + b3\text{Com}_t + \varepsilon_t$	1.03	0.69	15.0	8.3
$(\text{Palm}_t - \beta\text{Nasdaq}_t) = a + b(3\text{Com}_t - \beta\text{Nasdaq}_t) + \varepsilon_t$	0.95	0.54	13.4	10.0

Table 1. Regressions of Palm returns on 3Com returns. σ units are daily percent returns. Sample: 03 March 2000 - 27 July 2000.

Volatility and delinking. Table 1 demonstrates that something really weird (or really monetary) is going on in the Palm and 3Com valuation. All of 3Com’s value – and then some – is due to its Palm shares. The only way that 3Com and Palm returns can diverge is if there is news about the *rest* of 3Com. But when the stub is small, it would take *astounding* volatility of news about the rest of 3Com to account for the poor correlation between 3Com and Palm shares.

Here is the argument quantitatively. The value of 3Com should be (if the law of one price holds) the value of its components:

$$3\text{Com value} = \text{value Palm shares held by 3Com} + 3\text{Com stub value} \quad (1)$$

Equation (1) implies that the 3Com stub return can be inferred 3Com and Palm returns by

$$3\text{Com Stub Return}_t = \frac{(\text{Total 3Com value})_{t-1} 3\text{Com Return}_t}{(\text{3Com stub value})_{t-1}} - \frac{(\text{value of palm shares held by 3Com})_{t-1} \text{Palm Return}_t}{(\text{3Com stub value})_{t-1}}$$

As the stub value declines to zero, any deviation between Palm returns and 3Com returns implies infinite variation in the 3Com stub return.

Figure 6 shows the imputed 3Com stub return before the spinoff and the actual 3Com return after the spinoff, together with 15 day moving average standard deviations. It's clear that the "3Com stub risk" before the spinoff is much greater than the actual 3Com stub risk revealed after the spinoff. The overall standard deviation of the 3Com stub return is 32.2 percent (daily return!) before the spinoff and only 6.72 percent afterwards.

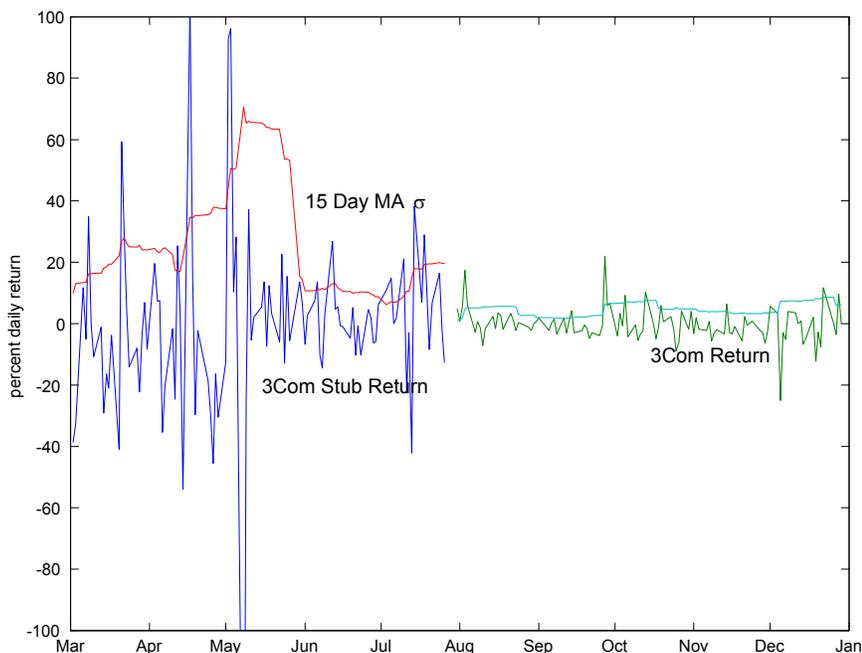


Figure 6: Return of the 3com stub (3Com less the value of Palm shares held by 3Com) before the spinoff and returns of 3Com itself after the spinoff.

At short horizons, Palm prices and 3Com prices are delinked. If you buy 3Com in order to bet on Palm shares, you are taking on much more than the true risks to the

stub value of 3Com. You are taking on the risk that when the exact moment comes to sell, the price move in the 3Com market will be delinked from the price move in the Palm market. *If you want to bet on Palm, you have to buy Palm.*

High Frequency Delinking. High frequency traders can't trade at the closing prices graphed in daily data, and they spend a lot of time looking at intraday prices, figuring out exactly when the right time is to buy and sell. Figure 7 shows a very close-up shot of one afternoon, March 14. (I chose the date at random among days other than the opening day, rather than select for a pretty plot. The opening day is completely chaotic.) The Figure presents the NYSE TAQ data of every single trade. There is a wide trading band at any point in time – good execution is important for day traders! The figure also shows substantial lack of correlation between 3Com and Palm. In this case 3Com is moving around in ways not followed by Palm. From the perspective of a high frequency trader, 3Com is a very poor substitute for Palm – even though all of 3Com's value and more is due to its holding of Palm shares.

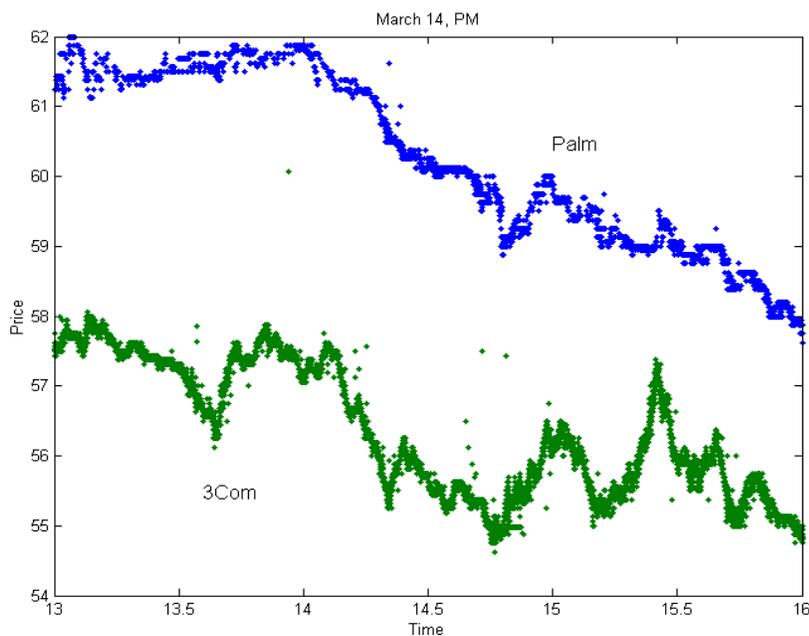


Figure 7: Trade by trade prices for 3Com and Palm on the afternoon of March 14 2000.

Options are delinked. What about options? Rather than buy Palm directly or buy 3Com, why not synthesize Palm with a long call and short put – or just buy the naked option corresponding to the direction in which you think Palm will go?

Lamont and Thaler (2001) look at options, and find massive violations of put-call parity. The synthetic 3Com stub is almost never negative. The synthetic Palm price is typically far below the actual Palm price, and does not allow an arbitrage. For

example, they report in their Table 6 that on March 17, the synthetic Palm was between 8 and 30 percent (depending on maturity) less than the actual Palm. They do not report a time series (and I have not yet constructed one), but if the *level* of synthetic Palm can be so far away from actual Palm, it seems likely that *changes* are also poorly correlated, in the same way that 3Com and Palm are poorly correlated. If you bet on Palm by buying much cheaper synthetic Palm, when the news you anticipate hits actual Palm it may not hit the synthetic Palm equally.

6. Why did 3Com Fall?

It’s interesting that on the day of the Palm IPO, 3Com *fell* 21%. One would think that the surprising good news about the value of Palm would *lift* 3Com, since 3Com holds 95% of Palm’s shares. Where is the bad news for 3Com (including 95% of Palm)?

A convenience yield view gives a reason for 3Com to fall. Until Palm is issued, the only way to bet on the prospects of Palm is to buy 3Com. The minute Palm starts trading, however, you can start to bet on Palm by buying and selling Palm directly. If the day-traders abandon 3Com and head off to trade Palm stock, 3Com loses much of its potential convenience yield, its required return rises, and its price falls. Of course, there must be some surprise in this shift in trading. Any perfectly predictable event should already be reflected in the price, leading to a slow relative decline in 3Com rather than a one-day 21% drop, as the perfectly predictable event of the actual spinoff did not seem to give a huge decline in Palm. Still, it’s possible that it was unclear whether the massive pre-IPO trading in 3Com would all move to Palm, given how few shares would be issued, and that there was some surprise in this event.

Summary

Table 2 summarizes many of the empirical similarities between the 3Com / Palm “mispricing” and the Dollar/Tbill “mispricing,” to show the many common features of the two phenomena.

	3Com/Palm	Dollar/Tbill
Law of one price violated	x	x
Restrictions on long-term short	x	x (no banknotes)
High turnover, short horizon in “expensive” end	x	x
Turnover higher as price spread higher	x	x (velocity)
High price security is “special” for trading	x	x
Price spread higher as quantity lower	x	x
Price spread lower as substitutes arise	x	x
Much shorting. substitutes despite cost	x	x (checking accounts)
Size can be large	x	x (hyperinflations)

Table 2. Summary of the similarities between 3Com/Palm and Dollar/Tbill mispricing

2.2 Making sense of a convenience yield for stocks

Does a trading-related convenience yield for stocks make any economic sense? If so, can it possibly be large enough to explain events such as 3Com / Palm or the Nasdaq?

2.2.1 Similar effects and the size of convenience prices

The size of monetary spreads. While we usually think of interest rates and convenience yields as small, the money/bond price spread has, historically, far exceeded the Palm/3Com spread. All it takes is a hyperinflation. For example, with a 100% inflation rate, the price of a one year bond is 1/2 (I assume a zero real rate to keep computations simple), so money is overpriced relative to one year bonds by 100%! And hyperinflations of many hundreds and thousands of percent have been common even in the 20th century. The Palm/3Com spread – arguably one of the largest convenience yields in a century of asset price data – is, at about a 50% per year relative return, mild by these standards.

It may seem strange to think of money as *overvalued* in a hyperinflation, but that is in fact the case – money suffers a high rate of depreciation, so its price today is “too high” relative to its price tomorrow. People hold a *small* amount of money, for a *short* time, because it is in fact extraordinarily *valuable* in a hyperinflation. I highlight the adjectives, as all three elements show up for Palm.

Convenience yields in finance Transactions-related price spreads and convenience yields in rates of return are not a novelty to finance. Of course, the commodity convenience yield is the heart of commodity futures pricing. Something about having grain in a warehouse, rather than relying on your ability to buy it if needed, must be valuable.

Krishnamurthy (2001) documents the spread between 30 year bonds, which are actively traded, and 29 year bonds, which are not. The 30 year bonds are “overpriced,” which is explainable by a similar transactions demand or convenience yield for holding them. As with Palm / 3Com, Krishnamurthy documents short constraints in the form of a “special repo rate.”

Krishnamurthy finds the overpricing is typically about 10 basis points. This is not as small as it seems. Ten basis points for a 30 year coupon bond is about \$1.50 per \$100 face value. In addition, the 29/30 year yield spread vanishes in a year, so the *expected return* spread is about 1.5 percentage points.

Mason (1987) and Boudoukh and Whitelaw (1991, 1993) documents the “benchmark” effect in Japanese bonds. Bond traders usually focus on a single issue called the “benchmark” which can trade as much as one percentage point above bonds with nearly identical terms. When the benchmark changes, the old benchmark declines in value and the new benchmark gains, in a manner reminiscent of 3Com’s fall and Palm’s rise when Palm opened. One percentage point on a 10 year (par, coupon, 5%

yield) bond is an \$8 price difference.

These numbers are smaller than Palm/3Com, but so is bond volatility. Furthermore, we expect that the differences in opinion or information about the value of government bonds is much smaller than that for a tech-stock IPO in March of 2000. Therefore, we expect smaller convenience yields.

Fernald and Rogers (2001) examine the relative prices of the two classes of shares of Chinese companies that can be held by domestic residents and foreigners, respectively. Though identical in other respects, the shares available to foreigners have traded for as little as *one-fourth* of the price of domestic shares. They also document that the domestic shares have as much as twice the price volatility as the foreign shares, and domestic shares average 2.4% daily turnover, compared to 0.3 % for foreign shares. A natural interpretation is that the domestic share market is where the high frequency trading takes place, so domestic shares inherit a convenience yield not given to the foreign shares. One-fourth seems like a lot, but Fernald and Rogers calculate that the price difference only translates to about a 4 percentage point per year difference in expected return, since stocks are long-lived securities.

The Gordon growth formula

$$\frac{p}{d} = \frac{1}{r - g}$$

reminds us that stocks whose prices are already high, due to high growth g or low required return r , are especially sensitive to changes in required return. For example, a stock with $r - g = 2\%$ has $p/d = 50$ while a stock with $r - g = 5\%$ has $p/d = 20$. A one percentage point reduction in required return raises the high priced stock to $p/d = 100$ while it only raises the low priced stock to $p/d = 25$. Based on this simple argument, it at least is not a surprise that convenience yield effects seem most dramatic among “growth stocks” whose prices are already high relative to book value, earnings or dividends.

3 Theories of a convenience yield

One’s first objection to the T bill/ Dollar bill analogy might be that money is special because it is the medium of exchange. You don’t use Palm stock to buy your morning coffee. However, the medium of exchange need not have a convenience yield, and a convenience yield is not limited to the medium of exchange. If every point of sale had a debt card reader, we could arrange our transactions without holding any money. Money would lose its convenience yield, even though it remained medium of exchange. Conversely, bond liquidity spreads, commodity convenience yields, and depreciating inventories of all kinds remain, though they are not media of exchange.

Money’s convenience yield comes because you have to hold some money as an inventory, or buffer stock. That feature is generated by the medium of exchange function when markets are not constantly available, but lots of other trading frictions

can generate an inventory demand.

Trading requires some shares of stock. Obviously, you can't buy stock if there is none. You can't short it either. To short stock, you have to find someone to borrow it from, and then you have to find someone to sell it to. You have to hold stock for a short time in order to short it. Duffie, Gârleanu and Pedersen (2001) model these short selling mechanics. In the model, everyone wants to short. An equilibrium short rebate keeps each agent happy to hold depreciating stock for the time it takes to find a borrower or a purchaser. The stock price falls slowly as short sales build up.

Scheinkman and Xiong (2002), following Harrison and Kreps (1978), present a model close to the verbal analysis I have presented. They show how a limited number of shares plus a large volume of trading based on differences of opinion can drive prices up above the frictionless "fundamental value."

So far, fully-rational models have not derived the extraordinary volume of trading. Traders say they act on differences of information, or differences of opinion (different processing of the same information). Yet every trader can't have better than average information. Scheinkman and Xiong rely on "overconfidence" – each trader thinks his signal is more accurate than it really is. Baker and Stein (2002) present a model based on irrational investors who underreact to news contained in order flow. Their model predicts that high turnover will be associated with low returns, i.e. high prices. Other models rely on "noise traders" or other less than fully rational behaviors.

Yet "irrationality" seems a superficial response to the deep problem posed by volume. Even hyper-rational finance professors and Nobel Prize winners run to high-frequency trading when they go off to Wall Street. Without high-frequency trading, the NYSE and NASDAQ would disappear, as retail life-cycle investing could not support them. Do we really believe that the NYSE and NASDAQ are temples to irrationality – that they would disappear if traders would only take to heart the fact that we can't all be better informed than average? Furthermore, when we write models that make arbitrary behaviors structural, we risk repeating the ill fate of Keynes' consumption function and Phillips' static curve.

I don't propose to solve the volume mystery here. It seems a productive first step to note that there is a lot of high frequency trading. Let us see trading can induce a convenience yield and inventory demand for shares sufficient to explain puzzling price patterns, and worry about micro foundations later. If irrational *trading*, but not irrational *valuation*, is behind puzzling price anomalies, that is already news.

It may give some comfort to remember that the microfoundations of money demand are also an enduring theoretical puzzle. Obviously, a dollar is helpful in arranging transactions, but vaguely realistic models of optimal transactions demand do not come close to the \$500 per capita of actual US cash holdings, and even less so the remaining demand for money in high inflation countries. The puzzles are the same – "why don't people use foreign currencies, or credit cards?" sounds a lot like "why didn't Palm traders use options, or 3Com?"

A basic theorem of frictionless finance (essentially, the Modigliani-Miller theorem) is that the value of the company should not depend on the number of shares issued to the public. Even if 99% of the shares are locked up, or held passively by some other company, the total market value should not be affected by the number of shares publicly issued. The convenience yield violates this theorem; it gives rise to a downward sloping “demand for shares,” especially in times of huge volume and much high frequency trading.

4 The Rise and Fall

The Palm/3-Com event is a microcosm of a larger puzzle—the Nasdaq rise and subsequent fall. Price volatility was large, share turnover enormous, short interest large, and many stocks, and especially many of the unusually highly priced stocks, also had the curious feature that only a small fraction of total shares were issued to the public. Did a convenience yield explain or at least contribute to the events? I review the same suggestive litany of facts in the larger experience.

The “bubble” was concentrated.

Figure 8 compares three indices, the NYSE, the NASDAQ and an index of NASDAQ tech stocks. (The tech index is composed of NASDAQ stocks with SIC code 737. I did not want an index composed of surviving stocks. 737 is the most common SIC code for stocks on various tech or internet stock indices. This is a narrow index – for example, it leaves out eBay – but for this purpose it is better to be too narrow than too wide.)

The figure makes clear that if there was a “bubble,” or some behavioral overenthusiasm for stocks, it was concentrated on Nasdaq stocks, and Nasdaq tech and internet stocks in particular.

Prices and volume are correlated in time series and cross section.

Figure 9 presents dollar volume for the three indices of Figure 8. The pattern is clear – both in the time series and in the cross-section, *prices are high where dollar volume is high*. (Ofek and Richardson 2001 Figures 1 and 2 present a narrower and even more dramatic internet index. Their index rises 10 times more than the S&P500, and shows the same strong correlation with volume as between the figures here.)

Dollar volume is the economically right measure. A high dollar volume will naturally result from a constant share volume and higher prices. That is still economically important higher volume, but it is interesting to verify whether this is the case. Figure 10 verifies that in fact share volume rose and fell as well as dollar prices of a given share volume.

The association of price and volume is a generic feature of the historical “bubbles,” though not of most theories applied to them. For example, Figure 11 shows a sur-

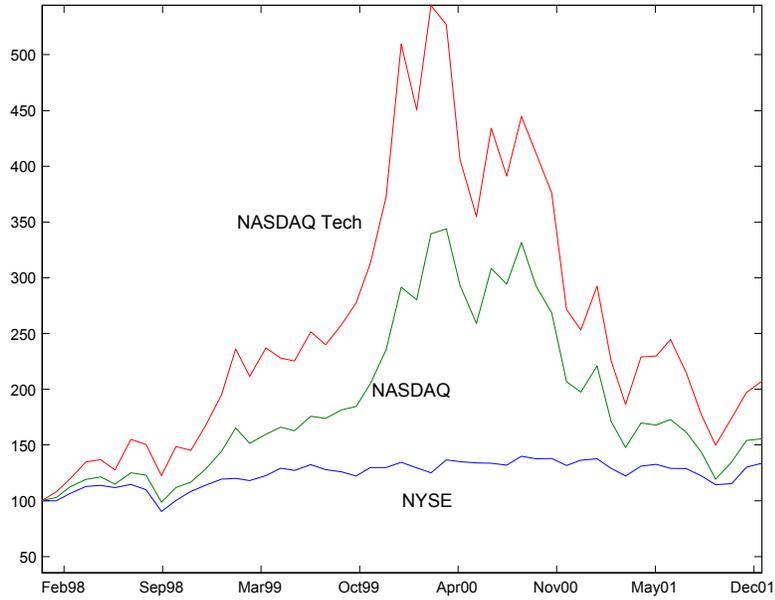


Figure 8: Total market value (price \times shares) of the NYSE, the NASDAQ, and NASDAQ stocks with SIC code 737, “computer processing, data processing.” Each index is based at 100 in Jan 1998. Source: CRSP.

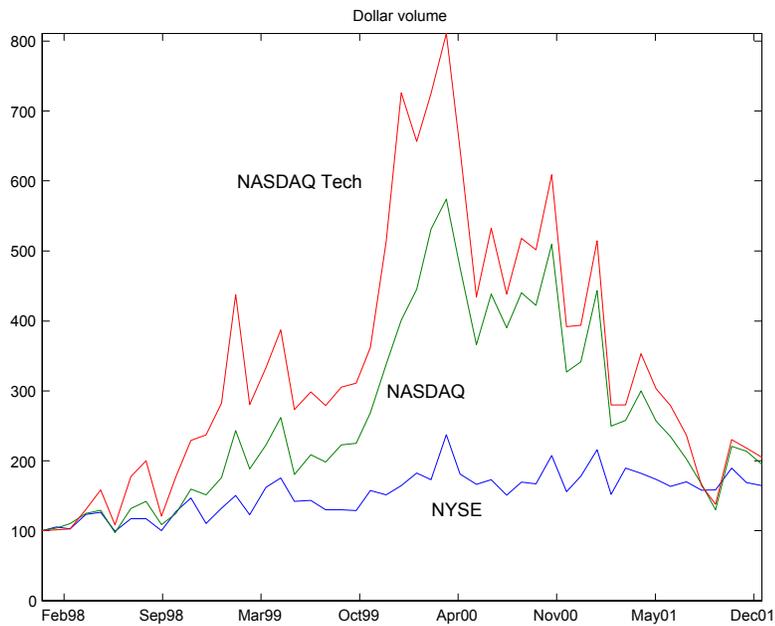


Figure 9: Dollar volume on NYSE, NASDAQ and NASDAQ with SIC code 737. Series are normalized to 100 on Jan 1 1998.

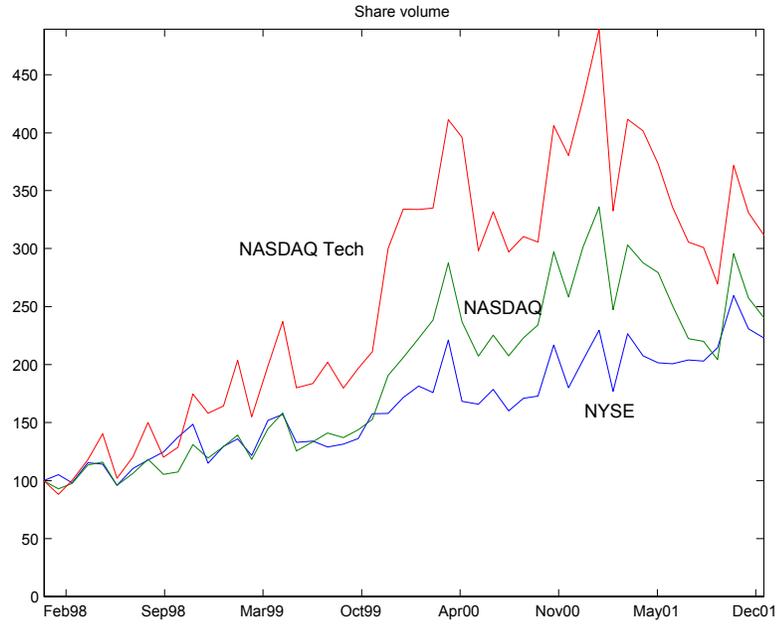


Figure 10: Share volume on NYSE, NASDAQ and NASDAQ with SIC code 737. Series are normalized to 100 in Jan 1998.

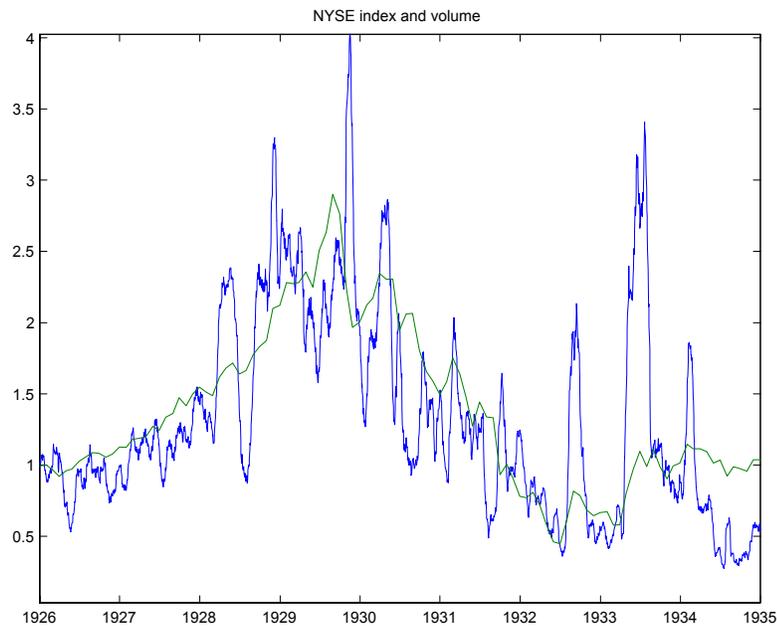


Figure 11: Share volume (20 day moving average) and NYSE index in the great crash. Volume is the more volatile series. NYSE index from CRSP; volume from NYSE. Both series normalized to 1 in 1926.

prising time-series correlation between the NYSE index and NYSE volume through the 1929 boom and crash. The graph presents *share* volume; *dollar* volume rose and declined even more calamitously. (Jones 2001 also notes the sharp decline in volume with the great crash. Jones shows that other measures of liquidity including bid/ask spreads also declined sharply.) Any value that shares had in high frequency trading in 1929 had evaporated by 1933. The South Sea Bubble and Tulip bubble similarly came and went with trading frenzies (Garber 2001).

Cross sectional relation between price and volume

There is also a cross-sectional relation between price and volume. To quantify this relationship, I run cross-sectional regressions of market equity / book equity on share turnover. Book equity is an easily available, though imperfect, proxy for the frictionless fundamental value. Table 3 presents the cross sectional regressions, and Figure 12 plots the cross-sectional correlation between log market/book and log turnover at each month. (For a single regression, the coefficient, R^2 and correlation all carry the same information.)

Sample	a	b	t	FM t	R^2	ρ
All CRSP Dec 1999	0.85	0.33	27		0.12	0.34
NASDAQ Dec 1999	0.89	0.38	23		0.13	0.37
All CRSP 1996-2000 (averages)	0.83	0.21	18	7.5	0.06	0.24
NASDAQ 1996-2000 (averages)	0.85	0.23	16	7.2	0.07	0.27

Table 3. Cross sectional regressions of market value / book value on share turnover.

$$\ln\left(\frac{ME}{BE_i}\right) = a + b \ln\left(\frac{\text{share volume}}{\text{shares outstanding}}\right) + \varepsilon_i$$

The first two rows present a single cross-sectional regression for December 1999. The second two rows present averages over cross sectional regressions run in every month from Jan 1996 to Dec 2000. t gives the OLS cross-sectional regression t statistic, and the average of the monthly t statistic in the second two rows. FM t presents the Fama MacBeth t statistic, calculated from the time-series standard deviation of the cross sectional regression coefficients b_t dividend by $\sqrt{5}$. ρ is the correlation between log market/book and log turnover. Book values follow Fama and French (2002).

Both the figure and the table show an important cross-sectional correlation between value and turnover, as a convenience yield predicts. Figure 12 shows interesting variation over time in the cross-sectional correlation. The correlation is highest during the boom year 1999. That is consistent with a positive correlation induced by a convenience yield sticking out more over the noise, in the year that convenience yield

was most important to prices. (The plot also shows year effects, due to the once per year changes in book value. I don't have a story why changes in book value should change the cross-sectional correlation so much.)

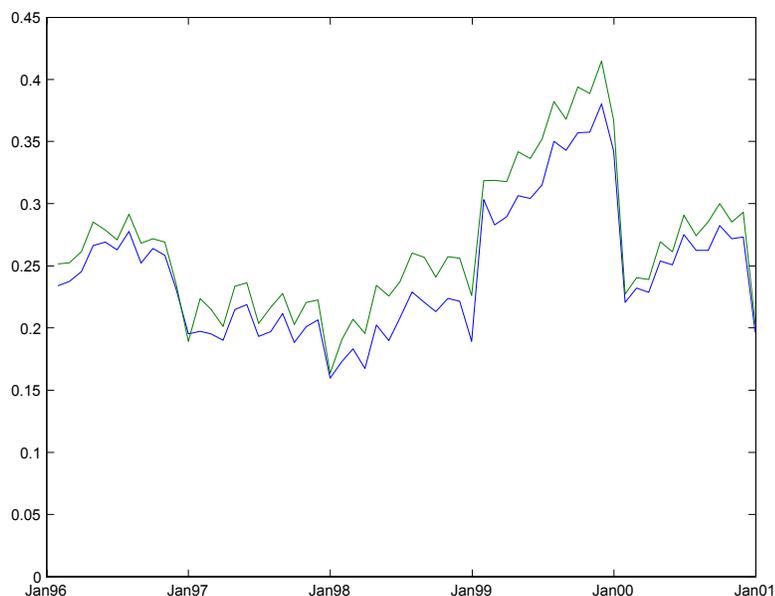


Figure 12: Cross-sectional correlations of log (market/book) with log turnover. Upper line: NASDAQ. Lower line: all CRSP stocks.

The regressions also suggest that the correlation is statistically significant. The OLS t statistics in the cross-sectional regressions are large, though the errors are undoubtedly cross-correlated. The Fama-MacBeth t statistics are also large, even given the short time dimension of the sample. (Since there is such a strong year effect visible in the plot, and as a rough correction for the time-series correlation of the cross-sectional regression coefficients, I formed Fama MacBeth standard errors by dividing the time series variance of the cross-sectional coefficients by 5, the number of years, rather than 60, the number of months.)

Other predictions for the internet bubble

Ofek and Richardson (2001) survey internet stocks in the late 1990s. In addition to the strong correlation between price and volume, they document all the other pieces of the convenience yield picture.

Short sales constraints. In addition to the usual problem that it is difficult and costly to maintain a long term short position for years, and that you may get wiped out in the mean time by arbitrage opportunities (Liu and Longstaff 2000), Ofek and Richardson document the exceptional difficulty of shorting internet stocks during the boom. Rebate rates were low, violations of put/call parity common, as was the case with Palm/3Com.

Small number of Shares available Most of the Ofek and Richardson internet index consists of new companies, and, as with Palm, small fractions of the total shares are initially available for trading. Insiders are typically “locked out” of selling for a fixed period, usually 180 days, after IPO.

Price declines when shares become available. One source of “share supply” is insider sales after lockup expiration. Many observers commented on the puzzle that prices fell when the lockup period ended, as if the market were assigning value to traded shares ignoring the presence of shares held by insiders. Ofek and Richardson (Figure 5) show that internet firms lost 12% of value (cumulative abnormal return) from 20 days prior to the end of lockup to 20 days after the end of lockup. They document a large fraction of shares removed from lockup between December 1999 and March 2000, coincident with the dot com peak (Figure 4) In addition, to the extent that traders want to bet on the internet as a whole rather than company-specific events, the huge supply of new shares via IPO counts to satisfy the transaction demand. The total supply of internet shares via IPO, SEO, and insider sales grew rapidly to a peak in March 2002, just as the price peaked as well (Figure 5).

5 A comparison of theories

What theories account for the curious behavior of tech stock prices? How many of the facts do they capture? Table 4 summarizes.

Fact	Theory					
	1	2	3	4	5	6
Prices rise, decline	x	x	x	x	x	x
Prices do not forecast earnings		x	x	x	x	x
Long term short difficult			x	x	x	x
Large dispersion of opinion					x	x
Price high with number shares low (ts & cs)					x	x
Price high with high volume (ts & cs)						x
Price high with volatility high (ts & cs)						x
Biggest in growth stocks						x

Table 4. Summary of theories and which facts they account for. 1: Frictionless model driven by earnings expectations 2: Frictionless model driven by risk premia. 3: Rational Bubble 4: Irrational valuation 5: Optimists’ opinions only with short restrictions. 6: Convenience yield, shares are needed for high frequency trading.

“ts&cs” stands for “time-series and cross-section.”

1. *Frictionless rational pricing – high earnings*

The standard model of frictionless pricing, for example the simple Gordon growth model with constant growth g and expected return r ,

$$\frac{p_t}{d_t} = \frac{1}{r - g}$$

can generate huge price rises if people expect huge earnings growth g . The internet was surely a once-in-a-generation technological novelty; a discovery on a par with the automobile or the computer itself. Couldn't prices have risen on rational expectations of earnings growth?

Many observers noted that the earnings growth required to support internet valuations was astronomically high. While fine on a story-telling, qualitative basis, this theory is difficult to match quantitatively. We now know that these observers were right. However, being *wrong*, once, is not the same thing as being *irrational*. Was it irrational to believe that the internet was a much bigger thing than it turned out to be?

Rational expectations can't be wrong all the time. If this time high prices were disappointed (so far) by ex-post earnings growth, at some other time, high prices must have been *underestimates* of even better earnings growth. In fact, the lessons of the volatility test literature starting with Shiller (1982) and the return forecasting literature such as Fama and French (1988), is that by and large high prices correspond to lower future returns, not to higher earnings growth.

In addition, the association of price movements with volume, scarce shares, short restrictions and so forth are just coincidences to this theory. That volume was high is as irrelevant as that the sky was blue.

2. *Frictionless rational pricing – low returns*

If not high g , perhaps low r is the explanation. If earnings grow at 5%, all it takes is a 5% required return to generate an *infinite* price. Absent an arbitrage opportunity, a time-varying risk premium can be invoked to explain any price pattern.

Fluctuations in the market as a whole might be sensibly understood as a time-varying risk premium. The top of the largest economic boom in postwar US history is exactly when you'd expect a risk premium to be low and stock prices to be high. This view is also completely consistent with the volatility test and return regression evidence – high prices *are* followed by low returns, and low returns can be perfectly rational. (Campbell and Cochrane 1999 is one model of this effect.)

Alas, while this story can easily explain procyclical variation in the market as a whole, it is much more strained to explain a boom in one particular segment of the market. Why should *internet* stocks suddenly have a much lower risk premium? This would have to come from a dramatic change not in overall risk aversion and the market premium, but from a dramatic change in the perceived covariance (beta) of internet earnings with respect to fundamental risk factors such as consumption, the

market portfolio, and so on. And then after March 2000, that beta changes back again. This is a hard story to swallow.

This model is also silent on a number of empirical regularities. Like all frictionless models, the huge volume of trading and the association with scarce shares are again irrelevant.

3. *Rational Bubbles*

A “rational bubble” occurs if everyone rationally holds stocks for any finite number of periods, but the “transversality condition” at infinity does not hold. For example, a stock with no dividend could satisfy

$$p_{t+1} = (1 + r)p_t + \varepsilon_{t+1}.$$

The expected return is $(1 + r)$, yet the price is not zero even though there are no dividends – the present value formula is violated. The price is expected to grow *forever*. Since we know the world will end in a few billion years, this model requires a small bit of irrationality – you’re counting on someone, someday, buying an overvalued stock. However, that irrationality is pushed arbitrarily far in the future.

Rational bubbles can produce fascinating price paths with booms and crashes. They also produce price volatility that does not forecast earnings *or* returns. Rational bubbles require a short constraint, otherwise you would short the bubble, pay any dividends and turn an arbitrage profit at infinity into an arbitrage profit now.

However, though “bubble” has been invoked to explain everything from tulip prices, the South Sea company, the 1929 Dow and 2000 tech stocks, this model is also silent on many common features of the events. Volume, again, is the most notable. In a “rational bubble” it is always rational to *hold* the stock for a period. There is no reason ever to buy or sell. It is also silent on all the other characteristics listed in Table 3. There is no reason the bubble should happen in tech stocks and not coal stocks, and no reason that “share supply” should seem to matter to prices.

4. *Irrational valuation*

Many observers including Lamont and Thaler conclude that this event is *prima facie* evidence for “irrational” investors, who simply put too high a value on Palm. The trouble with this explanation is that it can explain too much. If prices were too low, they would be labeled “irrational depression” rather than “irrational exuberance.” (And they have – most of the time the finance profession struggles with the equity premium puzzle that prices seem too low leading to too high average returns.) The strength of any theory is the number of things it *can’t* explain. I look forward to the day that behavioral finance can document structural patterns of behavior toward risk, well-understood evolutionary responses to our two million years of avoiding lions on the plains of Africa, that serve poorly in financial markets and can explain specific puzzles and not others. That day is not with us yet.

As with the other theories, irrational valuation is silent on many crucial attributes of the experience. Above all, volume: If investors were irrationally attracted to Palm

and bought too high, why did they get irrationally disgusted a day later and sell? Why did investors become irrationally attached to tech stocks and not coal stocks? Why did irrational attraction come and then go? Why did they lose their irrational attraction when insider's lockup periods expired?

One of the key definitions of “irrational” behavior, is that people recognize it when it's pointed out and change; education can overcome irrationality. This is one of the few ways to separate “irrationality” from just plain mistakes, or limited processing ability, and to give it some non-tautological content. The dog is not irrational for failing to short Palm; when the mistake is pointed out she just wags and pants as before. Yet the 3Com / Palm was widely noted in the business press (and at the GSB faculty lunchroom) on March 2, 2000; the overpricing was no secret. Similarly, story after story about how internet stocks were “overpriced” from the point of view of long-term investors appeared in the business press. People bought anyway, *but not for the long term* as the volume numbers show. Many finance academics were actively trading internet stocks – apparently, the best education in the world is not sufficient to overcome the desire to trade. Lamont and Thaler and Ofek and Richardson both document that institutional ownership was a little lower than usual in internet stocks, but “educated” institutions still held large amounts. Thus, we cannot explain Palm/3Com or the Nasdaq by educable irrationality.

5. *Short constraints mean we can't see the pessimists*

Ofek and Richardson (2001) interpret their findings as evidence for models in which short sales constraints keep the pessimists' views from being expressed, as in Miller (1977). Chen, Hong and Stein (2001) also advocate this view. This view also generates a “demand for shares.” In particular, as stressed by Ofek and Richardson, it explains the concentration of high prices in stocks with few shares outstanding. However, this view is again silent on volume.

Of course there is some blurring of categories here, especially given the number of models currently being explored. Models such as Harrison and Kreps (1978) have some trading; their speculative demand comes from the possibility that an even more optimistic investor will come along, and then today's optimist will sell to him. Rather than categorize papers, it's more productive to categorize the underlying stories. One story focuses on the fact that with short constraints we don't see the pessimists. This story does not fundamentally involve any volume; we can all wake up the next morning and see the same thing with no trades happening. If one adds changes in opinions or changes in traders, some volume emerges and, potentially, some even higher prices.

6. *Summary*

Table 4 summarizes the facts and theories. All of the theories explain prices that rise and then decline. Only the convenience yield theory says anything about volume. The convenience yield theory can *only* explain high prices if there is a restriction on the number of shares or substitutes and a large trading volume. All of the other theories are silent on this observation; high volume is a coincidence.

6 Concluding Remarks

A substantial convenience yield or liquidity premium for stocks that are undergoing a trading frenzy, based on wide dispersion in opinion about a fundamentally new technology, is an attractive explanation for the high prices of Nasdaq tech stocks around 2000. This view not only explains the rise and fall of prices, which many stories do, but it explains – indeed it *requires* – a wide variety of features common to this and similar experiences, including the association of high prices with high volume, high volatility, short sales frictions, and small numbers of available shares.

The possibility that some component of stock prices represents a trading-related convenience yield can help to explain a number of puzzles related to observations that demand curves seem to slope down at times. Among the most obvious, it suggests the long-run underperformance of IPOs (e.g. Loughran and Ritter 1995), following the Palm pattern on a more subdued scale. It suggests trading-related price anomalies such as the S&P 500 inclusion effect on prices (Harris and Gurel, 1986; Shleifer 1986) and betas (Barberis, Shleifer and Wurgler 2002): given the large number of index-linked contracts, stocks in the index will be traded a lot more than stocks out of the index.

The possibility of stock convenience yields is also interesting for policy and investment management. Obviously, buy-and hold investors and institutions should avoid stocks with high convenience yields, or lend them when short rebates make it profitable to do so. If convenience yield induced price distortions are substantial, fewer restrictions on trading and short-selling, and better substitutes such as individual stock futures markets can be important in limiting stock price volatility, and helping stocks to send the proper signals about physical investment rather than share issue.

7 References

- Barberis, Nicholas, Andrei Shleifer, and Jeffrey Wurgler, 2002, "Comovement," Manuscript.
- Baker, Malcom, and Jeremy C. Stein, 2002, "Market Liquidity as a Sentiment Indicator," Manuscript, Harvard University.
- Boudoukh, Jacob, and Robert Whitelaw, 1993, "Liquidity as a Choice Variable: A Lesson from the Japanese Bond Market," *Review of Financial Studies*, 6, 266-292.
- Boudoukh, Jacob, and Robert Whitelaw, 1991, "The Benchmark Effect in the Japanese Government Bond Market," *Journal of Fixed Income*, 1, 52-59.
- Brennan, Michael and Avanidhar Subrahmanyam, 1996, "Market Microstructure and Asset Pricing: On the Compensation for Illiquidity in Stock Returns," *Journal of Financial Economics* 41, 441-464.
- Campbell, John Y. and John H. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205-251.
- Chen Joseph, Harrison Hong and Jeremy C. Stein, 2001 "Breadth of Ownership and Stock Returns," NBER Working Paper 8151, forthcoming *Journal of Financial Economics*
- Duffie, Darrell, Nicolae Gârleanu and Lasse Heje Pedersen, 2001, "Valuation in Dynamic Bargaining Markets," Manuscript, Stanford University.
- Fama, Eugene F. and Kenneth R. French, 1988, "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics* 22: 3-27.
- Fama, Eugene F. and Kenneth R. French, 2002, "Testing Tradeoff and Pecking Order Predictions about Dividends and Debt," *Review of Financial Studies* 15, 1-33.
- Fernald, John and John H. Rogers, 2001, "Puzzles in the Chinese Stock Market," Forthcoming, *Review of Economics and Statistics*.
- Gallant, A. Ronald, Peter E. Rossi, George Tauchen, 1992, "Stock Prices and Volume," *The Review of Financial Studies*, 5, 199-242
- Harris, Lawrence and Eitan Gurel, 1986, "Price and Volume Effects Associated with Changes in the S&P 500: New Evidence for the Existence of Price Pressure," *Journal of Finance* 41, 851-860.
- Harrison, Michael, and David Kreps, 1978, "Speculative Investor Behavior in a Stock Market With Heterogeneous Expectations," *Quarterly Journal of Economics* 92, 323-336.

- Lamont, Owen, and Richard H. Thaler, 2001, "Can the Market Add and Subtract? Mispricing in Tech Stock Carve-outs," manuscript, University of Chicago
- Lamont, Owen, 2000, "Guilty as Charged: Violations of the Law of one Price in Financial Markets," Manuscript, University of Chicago.
- Liu, Jun and Francis Longstaff, 2000, "Losing Money on Arbitrages: Optimal Dynamic Portfolio Choice in Markets with Arbitrage Opportunities" Manuscript, UCLA.
- Loughran, Tim and Jay Ritter, 1995, "The New Issues Puzzle," *Journal of Finance*, 50, 23-51.
- Jones, Charles, 2001, "A Century of Stock Market Liquidity and Trading Costs," Manuscript, Columbia University.
- Jones, Charles and Owen Lamont, 2002, "Short sale constraints and stock returns," forthcoming, *Journal of Financial Economics*
- Krishnamurthy, Arvind, 2001, "The Bond/Old-Bond Spread," Manuscript, Northwestern University
- Ofek, Eli, and Matthew Richardson, 2001, "Dot Com Mania: The Rise and Fall of Internet Stock Prices" Manuscript, New York University,
- Ofek, Eli, and Matthew Richardson, 2002, "The IPO Lock-Up Period: Implications for Market Efficiency and Downward Sloping Demand Curves," manuscript, New York University
- Miller, Edward, 1977, "Risk, Uncertainty and Divergence of Opinion," *Journal of Finance* 32, 1151-1168.
- Mitchell, Mark, Todd Pulvino and Erik Stafford, 2002, "Limited Arbitrage in Equity Markets," *Journal of Finance* LVII, 551-584
- Scheinkman, José and Wei Xiong, 2002, "Overconfidence, Short-Sales Constraints and Bubbles" Manuscript, Princeton University.
- Shiller, Robert J., 1981, "Do Stock Prices Move too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review* 71, 421-436.
- Shleifer, Andrei, 1986, "Do Demand Curves for Stocks Slope Down?" *Journal of Finance* 41, 579, 590.