



American Finance Association

Mental Accounting, Loss Aversion, and Individual Stock Returns

Author(s): Nicholas Barberis and Ming Huang

Source: *The Journal of Finance*, Vol. 56, No. 4, Papers and Proceedings of the Sixty-First Annual Meeting of the American Finance Association, New Orleans, Louisiana, January 5-7, 2001 (Aug., 2001), pp. 1247-1292

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/2697796>

Accessed: 24/11/2009 12:02

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

Mental Accounting, Loss Aversion, and Individual Stock Returns

NICHOLAS BARBERIS and MING HUANG*

ABSTRACT

We study equilibrium firm-level stock returns in two economies: one in which investors are loss averse over the fluctuations of their stock portfolio, and another in which they are loss averse over the fluctuations of individual stocks that they own. Both approaches can shed light on empirical phenomena, but we find the second approach to be more successful: In that economy, the typical individual stock return has a high mean and excess volatility, and there is a large value premium in the cross section which can, to some extent, be captured by a commonly used multifactor model.

OVER THE PAST TWO DECADES, researchers analyzing the structure of individual stock returns have uncovered a wide range of phenomena, both in the time series and the cross section. In the time series, the returns of a typical individual stock have a high mean, are excessively volatile, and are slightly predictable using lagged variables. In the cross section, there is a substantial “value” premium, in that stocks with low ratios of price to fundamentals have higher average returns, and this premium can to some extent be captured by certain empirically motivated multifactor models.¹ These findings have attracted a good deal of attention from finance theorists. It has proved something of a challenge, though, to explain both the time series and cross-sectional effects in the context of an equilibrium model where investors maximize a clearly specified utility function.

In this paper, we argue that it may be possible to improve our understanding of firm-level stock returns by refining the way we model investor preferences. For guidance as to what kind of refinements might be important, we turn to the experimental evidence that has been accumulated on how people choose among risky gambles. Many of the studies in this literature suggest that *loss aversion* and *narrow framing* play an important

* University of Chicago Graduate School of Business and NBER, and Stanford University, Graduate School of Business, respectively. We are grateful to Michael Brennan, John Campbell, John Heaton, and Richard Thaler for discussing this paper at various conferences. We have also benefited from conversations with David Hirshleifer, Tano Santos, Andrei Shleifer, Jeremy Stein, and seminar participants at the University of Chicago, Harvard University, Yale University, and the NBER.

¹ The value premium was originally noted by Basu (1983) and Rosenberg, Reid, and Lanstein (1985); Fama and French (1992) provide more recent evidence. Fama and French (1993) show that a specific three-factor model can capture much of the value premium. Vuolteenaho (1999) documents the excess volatility and time series predictability of firm level stock returns.

role in determining attitudes towards risk. Financial economists do not typically incorporate these ideas into their models of asset prices. We investigate whether doing so can shed light on the behavior of individual stock returns.

Loss aversion is a feature of Kahneman and Tversky's (1979) descriptive model of decision making under risk, prospect theory, which uses experimental evidence to argue that people get utility from *gains and losses* in wealth, rather than from absolute levels. The specific finding known as loss aversion is that people are more sensitive to losses than to gains. Since our framework is intertemporal, we also make use of more recent evidence on dynamic aspects of loss aversion. This evidence suggests that the degree of loss aversion depends on *prior* gains and losses: A loss that comes after prior gains is less painful than usual, because it is cushioned by those earlier gains. On the other hand, a loss that comes after other losses is more painful than usual: After being burned by the first loss, people become more sensitive to additional setbacks.

A crucial question that arises in applying this evidence on loss aversion to the context of investing is: Over which gains and losses is the investor loss averse? Is he loss averse over changes in *total* wealth? Or is he loss averse over changes in the value of his portfolio of stocks or even over changes in the value of individual stocks that he owns? When gains and losses are taken to be changes in total wealth, we say that they are defined "broadly." When they refer to changes in the value of isolated components of wealth—the investor's stock portfolio or individual stocks that he owns—we say that they are defined "narrowly." Which gains and losses the investor pays attention to is a question about *mental accounting*, a term coined by Thaler (1980) to refer to the process by which people think about and evaluate their financial transactions.

Numerous experimental studies suggest that when doing their mental accounting, people engage in *narrow framing*, that is, they often appear to pay attention to narrowly defined gains and losses. This may reflect a concern for non-consumption sources of utility, such as regret, which are often more naturally experienced over narrowly framed gains and losses. If one of an investor's many stocks performs poorly, the investor may experience a sense of regret over the specific decision to buy that stock. In other words, individual stock gains and losses can be carriers of utility in their own right, and the investor may take this into account when making decisions.

In our analysis, we study the equilibrium behavior of firm-level stock returns when investors are loss averse and exhibit narrow framing in their mental accounting. We consider two kinds of narrow framing, one narrower than the other, and investigate whether either of them is helpful for understanding the data.

In the first economy we consider, investors get direct utility not only from consumption, but also from gains and losses in the value of individual stocks that they own. The evidence on loss aversion described above is applied to these narrowly defined gains and losses: The investor is loss averse over

individual stock fluctuations, and how painful a loss on a particular stock is depends on that stock's prior performance. We refer to this as "individual stock accounting."

In the second economy, investors get direct utility not only from consumption, but also from gains and losses in the value of their overall portfolio of stocks. The evidence on loss aversion is now applied to *these* gains and losses: The investor is loss averse over portfolio fluctuations, and how painful a drop in portfolio value is depends on the portfolio's prior performance. We call this "portfolio accounting," a form of narrow framing, although not as extreme as individual stock accounting.

In our first set of results, we show that for all its severity, individual stock accounting can be a helpful ingredient for understanding a wide range of empirical phenomena. In equilibrium, under this form of mental accounting, individual stock returns have a high mean, are more volatile than their underlying cash flows, and are slightly predictable in the time series. In the cross section, there is a large value premium: Stocks with low price-dividend ratios have higher average returns than stocks with high price-dividend ratios. Moreover, the same kinds of multifactor models that have been shown to capture the value premium in actual data can also do so in our simulated economy. At the same time, the model matches empirical features of *aggregate* asset returns. In equilibrium, aggregate stock returns have a high mean, excess volatility, and are moderately predictable in the time series, while the risk-free rate is constant and low.

Second, we find that the investor's system of mental accounting affects asset prices in a significant way. As we broaden the investor's decision frame from individual stock accounting to portfolio accounting, the equilibrium behavior of individual stock returns changes considerably: Their mean value falls, they become less volatile, and they also become more correlated with each other. Moreover, the value premium in the cross section disappears. Overall, portfolio accounting can explain some features of the data, but is less successful than individual stock accounting.

To understand where our results come from, consider first the case of individual stock accounting. Many of the effects here derive from a single source, namely a discount rate for individual stocks that changes as a function of the stock's past performance. If a stock has had good recent performance, the investor gets utility from this gain, and becomes less concerned about future losses on the stock because any losses will be cushioned by the prior gains. In effect, the investor perceives the stock to be less risky than before and discounts its future cash flows at a lower rate. Conversely, if one of his stocks performs dismally, he finds this painful and becomes more sensitive to the possibility of further losses on the stock. In effect, he views the stock as riskier than before and raises its discount rate.

This changing discount rate makes firm-level stock returns more volatile than underlying cash flows: A high cash flow pushes the stock price up, but this prior gain also lowers the discount rate on the stock, pushing the stock price still higher. It also generates a value premium in the cross section: In

this economy, a stock with a high price–dividend ratio (a growth stock) is often one that has done well in the past, accumulating prior gains for the investor, who then views it as less risky and requires a lower average return. A stock with a low price–dividend ratio (a value stock) has often had dismal prior performance, burning the investor, who now views it as riskier, and requires a higher average return. Finally, since the investor is loss averse over individual stock fluctuations, he dislikes the frequent losses that individual stocks often produce, and charges a high average return as compensation.

The reason the results are different under portfolio accounting is that in this case, changes in discount rates on stocks are driven by fluctuations in the value of the overall portfolio: When the portfolio does well, the investor is less concerned about losses on any of the stocks that he holds, since the prior portfolio gain will cushion any such losses. Effectively, he views *all* stocks as less risky. Discount rates on all stocks therefore go down simultaneously. Conversely, discount rates on all stocks go up after a prior portfolio loss.

This discount rate behavior is the key to many of the portfolio accounting results. Stock returns are less volatile here than under individual stock accounting. In the latter case, stocks are highly volatile because good cash-flow news is always accompanied by a lower discount rate, pushing the price up even more. Under portfolio accounting, good cash-flow news on a particular stock will only coincide with a lower discount rate on the stock if the portfolio as a whole does well. There is no guarantee of this, and so volatility is not amplified by as much. Since shocks to discount rates are perfectly correlated across stocks, individual stock returns are highly correlated with one another. Moreover, the value premium largely disappears since a stock's past performance no longer affects its discount rate, which is now determined at the portfolio level. Finally, while there is a substantial equity premium, it is not as large as under individual stock accounting. The investor is loss averse over portfolio level fluctuations, which are sizable but not as severe as the swings on individual stocks. The compensation for risk is therefore more moderate.

While individual stock accounting can potentially be a helpful way of thinking about the data, we emphasize that it is only a potential ingredient in an equilibrium model, and by no means a complete description of the facts. For one thing, we show that it underpredicts the correlation of stocks with each other, and argue that a model that combines individual stock accounting with broader forms of accounting is likely to be superior to a model that uses individual stock accounting alone.

The fact that we study equilibrium returns under both individual stock accounting and portfolio accounting is also useful for making additional predictions for future testing. If individual stock accounting is relatively more prevalent among individual investors as opposed to institutional investors, we would expect to see stocks held primarily by individuals—small stocks, for example—exhibit more of the features associated with individual stock

accounting. Other predictions arise, if, over time, investors change the way they do their mental accounting. For example, the increased availability of mutual funds since the early 1980s may have caused a shift away from individual stock accounting towards portfolio accounting, since funds automatically prevent investors from worrying about individual stock fluctuations. Our analysis predicts that stocks that were once held directly but are now held indirectly through mutual funds should exhibit specific changes in pricing behavior. Among other predictions, such stocks should have higher price-to-fundamentals ratios and exhibit a lower cross-sectional value premium.

Loss aversion and narrow framing have already been applied with some success to understanding the *aggregate* stock market. Benartzi and Thaler (1995) analyze the static portfolio problem of an investor who is loss averse over changes in his financial wealth and who is trying to allocate his wealth between T-bills and the stock market. They find that the investor is reluctant to allocate much to stocks, even if the expected return on the stock market is set equal to its high historical value. Motivated by this finding, Barberis, Huang, and Santos (2001) introduce loss aversion over financial wealth fluctuations into a dynamic equilibrium model and find that it captures a number of aggregate market phenomena. They do not address the time series or cross-sectional behavior of individual stocks. Moreover, since they consider only one risky asset, they cannot investigate the impact of different forms of mental accounting, which is our main focus in this paper.

Ours is not the only paper to address empirical phenomena like time series predictability and the cross-sectional value premium. Other promising approaches include models based on irrationality or bounded rationality, such as Barberis, Shleifer, and Vishny (1998), Daniel, Hirshleifer, and Subrahmanyam (2001), and Hong and Stein (1999); models based on learning, such as Brennan and Xia (2001); and models based on corporate growth options such as Berk, Green, and Naik (1999) and Gomes, Kogan, and Zhang (2001).

The rest of the paper is organized as follows. In Section I, we propose two different specifications for investor preferences: In one case, the investor is loss averse over fluctuations in the value of individual stocks in his portfolio; in the other case, he is loss averse only over fluctuations in overall portfolio value. Section II derives the conditions that govern equilibrium prices in economies with investors of each type. In Section III, we use simulated data to analyze equilibrium stock returns under each of the two kinds of mental accounting. Section IV discusses the results further and in particular, argues that they may be robust to generalizations that allow for heterogeneity across investors. Section V concludes.

I. Two Forms of Mental Accounting

Extensive experimental work suggests that *loss aversion* and *narrow framing* are important features of the way people evaluate risky gambles. In this section, we construct preferences that incorporate these two ideas.

Loss aversion is a central feature of Kahneman and Tversky's (1979) prospect theory, a descriptive model of decision making under risk, which argues that people derive utility from changes in wealth, rather than from absolute levels. The specific finding known as loss aversion is that people are more sensitive to reductions in wealth than to increases, in the sense that there is a kink in the utility function. A simple functional form that captures loss aversion is

$$w(X) = \begin{cases} X & \text{for } X \geq 0 \\ 2X & \text{for } X < 0 \end{cases}, \quad (1)$$

where X is the individual's gain or loss, and $w(X)$ is the utility of that gain or loss.

Kahneman and Tversky (1979) introduce loss aversion as a way of explaining why people tend to reject small-scale gambles of the form²

$$G = (110, \tfrac{1}{2}; -100, \tfrac{1}{2}).$$

Most utility functions used by financial economists are *not* able to explain these risk attitudes because they are differentiable everywhere, making the investor risk-neutral over small gambles.³

To incorporate loss aversion into an intertemporal framework, we need to take account of its dynamic aspects. Tversky and Kahneman (1981) note that their prospect theory was originally developed only for one-shot gambles and that any application to a dynamic context must await further evidence on how people think about sequences of gains and losses.

A number of papers have taken up this challenge, conducting experiments on how people evaluate sequences of gambles. In particular, Thaler and Johnson (1990) find that after a gain on a prior gamble, people are more risk seeking than usual, while after a prior loss, they become more risk averse. The result that risk aversion goes down after a prior gain, confirmed in other studies, has been labeled the "house money" effect, reflecting gamblers' increased willingness to bet when ahead.⁴ Thaler and Johnson (1990)

² This should be read as: "receive \$110 with probability $\frac{1}{2}$, and lose \$100 with probability $\frac{1}{2}$."

³ One exception is first-order risk aversion preferences, studied by Epstein and Zin (1990), Segal and Spivak (1990), Gul (1991), and others. However, this specification does not allow for narrow framing, which is central in our analysis. Of course, even if a utility function is differentiable, one can explain aversion to small-scale risks by increasing the function's curvature. However, this immediately runs into other difficulties. Rabin (2000) shows that if an increasing, concave, and differentiable utility function is calibrated so as to reject G at all wealth levels, then that utility function will also reject extremely attractive large-scale gambles, a troubling prediction.

⁴ It is important to distinguish Thaler and Johnson's (1990) evidence from other evidence presented by Kahneman and Tversky (1979) showing that people are risk averse over gains and risk seeking over losses; indeed this evidence motivates a feature of prospect theory that we do not consider here, namely the concavity (convexity) of the value function in the domain of gains

interpret these findings as evidence that the degree of loss aversion depends on prior gains and losses: A loss that comes after prior gains is less painful than usual, because it is cushioned by those earlier gains. A loss that comes after other losses, however, is more painful than usual: After being burned by the first loss, people become more sensitive to additional setbacks.

An important modeling issue that remains is: Which gains and losses should this evidence on loss aversion be applied to? Are investors only loss averse over “broadly” defined gains and losses, by which we mean changes in *total* wealth, or are they also loss averse over “narrowly” defined gains and losses, namely changes in the value of isolated components of their wealth, such as their stock portfolio or even individual stocks that they own? Which gains and losses people focus on is a question about what Thaler (1980) calls *mental accounting*, the process people use to think about and evaluate their financial transactions.

To see why mental accounting matters, consider the following simple example. An investor is thinking about buying a portfolio of two stocks—one share of each, say. The shares of both stocks are currently trading at \$100, and after careful thought, the investor decides that for both stocks, the share value a year from now will be distributed as

$$(150, \tfrac{1}{2}; 70, \tfrac{1}{2}), \quad (1)$$

independently across the two stocks.

Suppose that the investor’s loss aversion is captured by the functional form in equation (1). If he is loss averse over portfolio fluctuations, the expected utility of the investment is⁵

$$\tfrac{1}{4}w(100) + \tfrac{1}{2}w(20) + \tfrac{1}{4}w(-60) = 5,$$

while if he is loss averse over individual stock fluctuations, it is⁶

$$2[\tfrac{1}{2}w(50) + \tfrac{1}{2}w(-30)] = -10,$$

which is not as attractive.

(losses). One set of evidence pertains to one-shot gambles, the other to sequences of gambles. Kahneman and Tversky’s (1979) evidence suggests that people are willing to take risks in order to avoid a loss; Thaler and Johnson’s (1990) evidence suggests that if these efforts are unsuccessful and the investor suffers an unpleasant loss, he will *subsequently* act in a more risk-averse manner.

⁵ This calculation says: with probability $\frac{1}{4}$, both stocks will gain \$50, for a total gain of \$100; with probability $\frac{1}{2}$, one stock will gain \$50, the other will lose \$30, for a total gain of \$20; and with probability $\frac{1}{4}$, both stocks will lose \$30, for a total loss of \$60.

⁶ This calculation says: for each stock, there is an equal chance of a gain of \$50 and a loss of \$30.

Which form of mental accounting is a better description of individual behavior? Traditional asset-pricing models usually assume as broad a form of accounting as possible: Utility is typically specified only over *total* wealth or over consumption, and not over individual stock fluctuations. A substantial body of experimental work, however, suggests that when doing their mental accounting, people engage in *narrow framing*, that is, they often do appear to focus on narrowly defined gains and losses.

The studies in this literature typically ask subjects who are already facing some kind of financial risk whether they would be willing to accept an additional, uncorrelated gamble. This gamble is designed to be unattractive when viewed in isolation: anyone focusing on the gains and losses of the gamble itself would tend to reject it. On the other hand, the gamble is designed to be attractive when viewed as part of the subjects' broader portfolio of gambles: subjects who are able to focus on portfolio gains and losses would tend to accept the gamble. Typically, these studies find that the additional gamble is rejected, suggesting that subjects do frame narrowly to some extent.⁷

The absence of narrow framing from standard asset-pricing models is probably due to doubts about its normative acceptability. These doubts may be unwarranted: Narrow framing *can* be defended on normative grounds because it may simply reflect a concern for non-consumption sources of utility, which are often naturally experienced over narrowly defined gains and losses. Regret is one example of such utility: A loss is more painful to us if it is linked to an action we took than if it simply befalls us through no fault of our own. If one of an investor's many stocks performs poorly, the investor may experience a sense of regret over the specific decision to buy that stock. Since each stock is associated with a distinct decision, namely the decision to buy that particular stock, each stock's gains and losses can give rise to a distinct source of utility, based on regret or euphoria about the initial buying decision.

In other situations, narrow framing is less acceptable from a normative perspective. These are situations where it arises because of cognitive limitations: Even though we know that gains and losses in *total* wealth are more relevant for our consumption decisions, we may focus too much on gains and losses in one part of our wealth—in our stock portfolio—simply because information about those gains and losses is more readily available.

In what follows, we study asset prices in economies where investors are loss averse and exhibit narrow framing in their mental accounting. In the first economy we consider, investors get direct utility from consumption, but also from gains and losses in the value of individual stocks that they own. The evidence on loss aversion is applied to these narrowly defined gains and

⁷ Redelmeier and Tversky (1992), Kahneman and Lovallo (1993), Gneezy and Potters (1997), Thaler et al. (1997), Benartzi and Thaler (1999), and Rabin and Thaler (2000) present evidence of various kinds of narrow framing. Read, Loewenstein, and Rabin (1999) review some of the evidence and discuss possible explanations of why people frame decisions the way they do.

losses: The investor is loss averse over individual stock fluctuations, and how painful a loss on a particular stock is depends on that stock's prior performance. We refer to this as "individual stock accounting."⁸

In our second piece of analysis, we consider an economy where investors get direct utility from consumption, but also from gains and losses in the value of their overall portfolio of stocks. The evidence on loss aversion is now applied to *these* gains and losses: The investor is loss averse over portfolio fluctuations, and how painful a drop in portfolio value is depends on the portfolio's prior performance. We call this "portfolio accounting." While this is a broader form of mental accounting than individual stock accounting, it still represents narrow framing: The investor is segregating his stock portfolio from his other forms of wealth such as human capital, and is focusing on its fluctuations separately.

It is worth noting one alternative interpretation of individual stock accounting. Until the advent of mutual funds, individuals with holdings in the stock market typically owned only a very small number of stocks. As a result, even if these investors engaged in *portfolio* accounting—a weak form of narrow framing—the effect on prices may be more accurately captured by our model of *individual stock* accounting. In other words, individual stock accounting may be primarily a consequence of the highly undiversified portfolios held by many investors, and not necessarily the result of excessively narrow framing.

We now show how these two forms of mental accounting can be incorporated into a traditional asset-pricing framework, starting with individual stock accounting in Section I.A and then moving to portfolio accounting in Section I.B. In both cases, there are two kinds of assets: a risk-free asset in zero net supply, paying a gross interest rate of $R_{f,t}$ between time t and $t + 1$; and n risky assets—"stocks"—each with a total supply of one unit. The gross return on stock i between time t and $t + 1$ is $R_{i,t+1}$.

A. Individual Stock Accounting

When the investor is loss averse over individual stock fluctuations, he chooses consumption C_t and an allocation $S_{i,t}$ to stock i to maximize

$$E \sum_{t=0}^{\infty} \left[\rho^t \frac{C_t^{1-\gamma}}{1-\gamma} + b_0 \bar{C}_t^{-\gamma} \rho^{t+1} \sum_{i=1}^n v(X_{i,t+1}, S_{i,t}, z_{i,t}) \right]. \quad (2)$$

⁸ A skeptic could argue that an investor who does individual stock accounting will be reluctant to take on blatantly attractive opportunities, such as exploiting a relative mispricing between two stocks by going long one and short the other. Even if he is *sure* to make \$5 on the long position and to lose only \$3 on the short, he may code this as $5 - 2(3)$, which does not look attractive. However, since the long and short positions are really components of a single trading idea, it is more likely that the investor will evaluate the strategy as a single entity: he will code a gain of $5 - 3 = 2$, and will be keen to take on the opportunity.

The first term in this preference specification, utility over consumption C_t , is a standard feature of asset-pricing models. Although the framework does not require it, we specialize to power utility, the benchmark case studied in the literature. The parameter ρ is the time discount factor, and $\gamma > 0$ controls the curvature of utility over consumption.⁹

The second term models the idea that the investor is loss averse over changes in the value of individual stocks that he owns. The variable $X_{i,t+1}$ measures the gain or loss on stock i between time t and time $t + 1$, a positive value indicating a gain and a negative value, a loss. The utility the investor receives from this gain or loss is given by the function v , and it is added up across all stocks owned by the investor. It is a function not only of the gain or loss itself, but also of $S_{i,t}$, the value of the investor's holdings of stock i at time t , and of a state variable $z_{i,t}$, which measures the investor's gains or losses on the stock prior to time t as a fraction of $S_{i,t}$. By including $S_{i,t}$ and $z_{i,t}$ as arguments of v , we allow the investor's prior investment performance to affect the way subsequent losses are experienced.

As discussed earlier, we think of v as capturing utility unrelated to consumption. Regret is one example of this, but there may also be other kinds of non-consumption utility at work here. An investor may interpret a big loss on a stock as a sign that he is a second-rate investor, thus dealing his ego a painful blow, and he may feel humiliation in front of friends and family when word about the failed investment leaks out.

The $b_0 \bar{C}_t^{-\gamma}$ coefficient on the loss aversion terms is a scaling factor that ensures that risk premia in the economy remain stationary even as aggregate wealth increases over time. It involves per capita consumption $\bar{C}_t$ that is exogenous to the investor, and so does not affect the intuition of the model. The constant b_0 controls the importance of the loss aversion terms in the investor's preferences; setting $b_0 = 0$ reduces our framework to the much studied consumption-based model with power utility.

Barberis et al. (2001) have already formalized the notion of loss aversion in a model of the aggregate stock market. We borrow their specification, which we summarize in the remainder of this section. We limit ourselves to describing the essential structure; Barberis et al. (2001) provide more supporting detail.

The gain or loss on stock i between time t and $t + 1$ is measured as

$$X_{i,t+1} = S_{i,t} R_{i,t+1} - S_{i,t} R_{f,t}. \quad (3)$$

In words, the gain is the value of stock i at time $t + 1$ minus its value at time t multiplied by the risk-free rate. Multiplying by the risk-free rate models the idea that investors may only view the return on a stock as a gain if it exceeds the risk-free rate. The unit of time is a year, so that gains and losses

⁹ For $\gamma = 1$, we replace $C_t^{1-\gamma}/(1 - \gamma)$ with $\log C_t$.

are measured annually. While the investor may check his holdings much more often than that, even several times a day, we assume that it is only once a year, perhaps at tax time, that he confronts his past performance in a serious way.

The variable $z_{i,t}$ tracks prior gains and losses on stock i . It is the ratio of another variable, $Z_{i,t}$, to $S_{i,t}$, so that $z_{i,t} = Z_{i,t}/S_{i,t}$. Barberis et al. (2001) call $Z_{i,t}$ the “historical benchmark level” for stock i , to be thought of as the investor’s memory of an earlier price level at which the stock used to trade. When $S_{i,t} > Z_{i,t}$, or $z_{i,t} < 1$, the stock price today is higher than what the investor remembers it to be, making him feel as though he has accumulated prior gains on the stock, to the tune of $S_{i,t} - Z_{i,t}$. When $S_{i,t} < Z_{i,t}$, or $z_{i,t} > 1$, the current stock price is lower than it used to be, so that the investor feels that he has had past losses, again of $S_{i,t} - Z_{i,t}$.

The point of introducing $z_{i,t}$ is to allow v to capture experimental evidence suggesting that the pain of a loss depends on prior outcomes. This is done by defining v in the following way. When $z_{i,t} = 1$,

$$v(X_{i,t+1}, S_{i,t}, 1) = \begin{cases} X_{i,t+1} & \text{for } X_{i,t+1} \geq 0 \\ \lambda X_{i,t+1} & \text{for } X_{i,t+1} < 0 \end{cases}, \quad (4)$$

with $\lambda > 1$. For $z_{i,t} < 1$,

$$v(X_{i,t+1}, S_{i,t}, z_{i,t}) = \begin{cases} S_{i,t} R_{i,t+1} - S_{i,t} R_{f,t} & \text{for } R_{i,t+1} \geq z_{i,t} R_{f,t} \\ S_{i,t}(z_{i,t} R_{f,t} - R_{f,t}) + \lambda S_{i,t}(R_{i,t+1} - z_{i,t} R_{f,t}) & \text{for } R_{i,t+1} < z_{i,t} R_{f,t} \end{cases}, \quad (5)$$

and for $z_{i,t} > 1$,

$$v(X_{i,t+1}, S_{i,t}, z_{i,t}) = \begin{cases} X_{i,t+1} & \text{for } X_{i,t+1} \geq 0 \\ \lambda(z_{i,t}) X_{i,t+1} & \text{for } X_{i,t+1} < 0 \end{cases}, \quad (6)$$

with

$$\lambda(z_{i,t}) = \lambda + k(z_{i,t} - 1), \quad (7)$$

and $k > 0$.

It is easiest to understand these equations graphically. Figure 1 shows the form of v : the solid line for $z_{i,t} = 1$, the dash-dot line for $z_{i,t} < 1$, and the dashed line for $z_{i,t} > 1$. When $z_{i,t} = 1$, the case where the investor has neither prior gains nor prior losses on stock i , v is a piecewise linear function




Figure 1. Utility of gains and losses. The utility function is shown for each of three cases: when the investor has prior gains (dash-dot line), prior losses (dashed line), or neither prior gains nor prior losses (solid line).

with a slope of one in the positive domain and a slope of $\lambda > 1$ in the negative domain. This gives it a kink at the origin where the gain equals zero and provides a simple representation of loss aversion.

When $z_{i,t} < 1$, the investor has accumulated prior gains on stock i . The form of $v(X_{i,t+1}, S_{i,t}, z_{i,t})$ is the same as for $v(X_{i,t+1}, S_{i,t}, 1)$ except that the kink is no longer at the origin but a little to the left; how far to the left depends on the size of the prior gain. This formulation captures the idea that prior gains may cushion subsequent losses. In particular, the graph shows that a small loss on stock i is penalized at the gentle rate of one, rather than λ : Since this loss is cushioned by the prior gain, it is less painful. If the loss is so large that it depletes the investor's entire reserve of prior gains, it is once again penalized at the more severe rate of $\lambda > 1$.

The remaining case is $z_{i,t} > 1$, where stock i has been losing value. The form of $v(X_{i,t+1}, S_{i,t}, z_{i,t})$ in this case has a kink at the origin just like $v(X_{i,t+1}, S_{i,t}, 1)$, but differs from $v(X_{i,t+1}, S_{i,t}, 1)$ in that losses are penalized at a rate more severe than λ , capturing the idea that losses that come after other losses are more painful than usual. How much higher than λ the penalty is is determined by equation (7), and in particular by the constant k . Note that the larger the prior loss, as measured by $z_{i,t}$, the more painful any subsequent losses will be.

To complete the model description, we need an equation for the dynamics of $z_{i,t}$. Based on Barberis et al. (2001), we use

$$z_{i,t+1} = \eta \left(z_{i,t} \frac{\bar{R}_i}{R_{i,t+1}} \right) + (1 - \eta)(1), \quad (8)$$

where $\bar{R}_i$ is a fixed parameter and $\eta \approx 1$. Note that if the return on stock i is particularly good, so that $R_{i,t+1} > \bar{R}_i$, the state variable $z_{i,t} = Z_{i,t}/S_{i,t}$ falls in value. This means that the benchmark level $Z_{i,t}$ rises less than the stock price $S_{i,t}$, increasing the investor's reserve of prior gains. In other words, equation (8) captures the idea that a particularly good return should increase the amount of prior gains the investor feels he has accumulated on the stock. It also says that a particularly poor return depletes the investor's prior gains: If $R_{i,t+1} < \bar{R}_i$, then $z_{i,t}$ goes up, showing that $Z_{i,t}$ falls less than $S_{i,t}$, decreasing $S_{i,t} - Z_{i,t}$. The parameter η controls the persistence of the state variable and hence how long prior gains and losses affect the investor. If $\eta \approx 1$, a prior loss, say, will increase the investor's sensitivity to further losses for many subsequent periods.

Implicit in equation (8) is an assumption that the evolution of $z_{i,t}$ is unaffected by any actions the investor might take, such as buying or selling shares of the stock. In many cases, this is reasonable: If the investor sells some shares for consumption purposes, it is plausible that any prior gains on the stock are reduced in proportion to the amount sold—in other words, that $z_{i,t}$ remains constant. More extreme transactions, such as selling one's entire holdings of the stock, might plausibly affect the way $z_{i,t}$ evolves. In assuming that they do not, we are making a strong assumption, but one that is very useful in keeping our analysis tractable.¹⁰

The parameter $\bar{R}_i$ is not a free parameter, but is determined endogeneously by imposing the requirement that in equilibrium, the median value of $z_{i,t}$ be equal to one. The idea behind this is that half the time, the investor should feel as though he has prior gains, and the rest of the time as though he has prior losses. It turns out that $\bar{R}_i$ is typically of similar magnitude to the average stock return.

B. Portfolio Accounting

The second form of narrow framing we consider is portfolio accounting, where investors are loss averse only over portfolio fluctuations. In particular, they choose consumption C_t and an allocation $S_{i,t}$ to stock i to maximize

$$E \sum_{t=0}^{\infty} \left[\rho^t \frac{C_t^{1-\gamma}}{1-\gamma} + b_0 \bar{C}_t^{-\gamma} \rho^{t+1} v(X_{t+1}, S_t, z_t) \right]. \quad (9)$$

¹⁰ An alternative way of interpreting this implicit assumption in equation (8) is that it represents a form of bounded rationality: When making his investment decisions, the investor is simply unable to figure out the effect of his actions on the future evolution of the state variable.

Here, X_{t+1} is the gain or loss on the investor's overall *portfolio* of risky assets between time t and time $t + 1$, $S_t = \sum_{i=1}^n S_{i,t}$ is the value of those holdings at time t , and z_t is a variable that measures prior gains and losses on the portfolio as a fraction of S_t . Once again, we interpret v as a non-consumption source of utility, which in this case is experienced over changes in overall portfolio value and not over changes in individual stock value.

Portfolio gains and losses are measured as

$$X_{t+1} = S_t R_{t+1} - S_t R_{f,t}, \quad (10)$$

where R_{t+1} is the gross return on the portfolio. When $z_t = 1$, v is defined as

$$v(X_{t+1}, S_t, 1) = \begin{cases} X_{t+1} & \text{for } X_{t+1} \geq 0 \\ \lambda X_{t+1} & \text{for } X_{t+1} < 0 \end{cases}, \quad (11)$$

with $\lambda > 1$. For $z_t < 1$,

$$v(X_{t+1}, S_t, z_t) = \begin{cases} S_t R_{t+1} - S_t R_{f,t} & \text{for } R_{t+1} \geq z_t R_{f,t} \\ S_t(z_t R_{f,t} - R_{f,t}) + \lambda S_t(R_{t+1} - z_t R_{f,t}) & \text{for } R_{t+1} < z_t R_{f,t} \end{cases}, \quad (12)$$

and for $z_t > 1$,

$$v(X_{t+1}, S_t, z_t) = \begin{cases} X_{t+1} & \text{for } X_{t+1} \geq 0 \\ \lambda(z_t) X_{t+1} & \text{for } X_{t+1} < 0 \end{cases}, \quad (13)$$

with

$$\lambda(z_t) = \lambda + k(z_t - 1), \quad (14)$$

and $k > 0$. Finally, the dynamics of z_t are given by

$$z_{t+1} = \eta \left(z_t \frac{\bar{R}}{R_{t+1}} \right) + (1 - \eta)(1). \quad (15)$$

In summary, the functional forms are identical to what they were in the case of individual stock accounting. The only difference is that in equation (2), the investor experiences loss aversion over changes in the value of each stock that he owns, while in equation (9), he is loss averse only over overall portfolio fluctuations.

II. Equilibrium Prices

We now derive the conditions that govern equilibrium prices in two different economies. The first economy is populated by investors who do individual stock accounting and have the preferences laid out in equations (2) through (8). Investors in the second economy do portfolio accounting, and have the preferences in equations (9) through (15). In both cases, there is a continuum of investors, with a total “mass” of one.

In each economy, we want to compute the price $P_{i,t}$ of stock i , say, which we model as a claim to a stream of perishable output given by the dividend sequence $\{D_{i,t}\}$, where

$$\log\left(\frac{D_{i,t+1}}{D_{i,t}}\right) = g_i + \sigma_i \varepsilon_{i,t+1}, \quad (16)$$

with $\varepsilon_t = (\varepsilon_{1,t}, \dots, \varepsilon_{n,t}) \sim \text{i.i.d. } N(0, \Omega)$, and where $\Omega_{ij} = (\omega_{ij})$ with $\omega_{ii} = 1$.

Aggregate consumption evolves according to

$$\log\left(\frac{\bar{C}_{t+1}}{\bar{C}_t}\right) = g_c + \sigma_c \eta_{t+1}, \quad (17)$$

where $\eta_t \sim N(0,1)$, i.i.d. over time, and

$$\text{corr}(\eta_t, \varepsilon_{i,t'}) = \begin{cases} \omega_{ci} & \text{for } t = t' \\ 0 & \text{for } t \neq t' \end{cases}. \quad (18)$$

We do not impose the Lucas (1978) restriction that aggregate consumption equals the aggregate dividend. The advantage of this is that it allows the volatility of consumption growth and of dividend growth to be very different in our model, as they are in the data. Given that aggregate consumption differs from the aggregate dividend, we fill the gap by assuming that each agent also receives a stream of nonfinancial income $\{Y_t\}$ —labor income, say. We assume that $\{Y_t\}$ and $\{D_{i,t}\}_{i=1,\dots,n}$ form a joint Markov process whose distribution gives $\bar{C}_t \equiv \sum_{i=1}^n D_{i,t} + Y_t$ the distribution in equation (17). For simplicity, we assume that agents are not loss averse over labor income fluctuations, although this can be relaxed without affecting the main features of our results.

A. Equilibrium Prices Under Individual Stock Accounting

Consider first an economy where investors have the preferences given in equations (2) through (8). Our assumptions so far allow us to construct a Markov equilibrium in which the risk-free rate is constant and the state variable $z_{i,t}$ determines the distribution of returns on stock i . Specifically,

we assume that the price–dividend ratio of stock i is a function of the state variable $z_{i,t}$,

$$f_{i,t} \equiv P_{i,t}/D_{i,t} = f_i(z_{i,t}), \quad (19)$$

and then look for an equilibrium satisfying this assumption. Under this one-factor assumption, the distribution of the stock return $R_{i,t+1}$ is determined by $z_{i,t}$ and the function $f_i(\cdot)$ as follows:

$$\begin{aligned} R_{i,t+1} &= \frac{P_{i,t+1} + D_{i,t+1}}{P_{i,t}} = \frac{1 + P_{i,t+1}/D_{i,t+1}}{P_{i,t}/D_{i,t}} \frac{D_{i,t+1}}{D_{i,t}} \\ &= \frac{1 + f_i(z_{i,t+1})}{f_i(z_{i,t})} \frac{D_{i,t+1}}{D_{i,t}} = \frac{1 + f_i(z_{i,t+1})}{f_i(z_{i,t})} e^{g_i + \sigma_i \varepsilon_{i,t+1}}. \end{aligned} \quad (20)$$

Intuitively, the value of stock i can change because of news about dividends $\varepsilon_{i,t+1}$, or because its price–dividend ratio $f_{i,t}$ changes. Changes in this ratio are driven by changes in $z_{i,t}$, which tracks the past performance of the stock. Recent gains (losses) on the stock make the investor perceive the stock as less (more) risky, changing its price–dividend ratio.

In equilibrium, and under rational expectations about stock returns and aggregate consumption levels, the agents in our economy must find it optimal to hold the market supply of zero units of the risk-free asset and one unit of each stock at all times, and to consume their labor income and the dividend on each stock every period.¹¹ The proposition below characterizes the equilibrium.¹²

PROPOSITION 1: *For the preferences in equations (2) through (8), necessary and sufficient conditions for a one-factor Markov equilibrium are*

$$R_f = \rho^{-1} e^{\gamma g_c - \gamma^2 \sigma_c^2 / 2}, \quad (21)$$

and

$$\begin{aligned} 1 &= \rho e^{g_i - \gamma g_c + (1/2) \gamma^2 \sigma_c^2 (1 - \omega_{ci}^2)} E_t \left[\frac{1 + f_i(z_{i,t+1})}{f_i(z_{i,t})} e^{(\sigma_i - \gamma \omega_{ci} \sigma_c) \varepsilon_{i,t+1}} \right] \\ &\quad + b_0 \rho E_t \left[\hat{v} \left(\frac{1 + f_i(z_{i,t+1})}{f_i(z_{i,t})} e^{g_i + \sigma_i \varepsilon_{i,t+1}}, z_{i,t} \right) \right], \end{aligned} \quad (22)$$

¹¹ We need to impose rational expectations about aggregate consumption because the agent's utility includes aggregate consumption as a scaling term.

¹² Throughout the paper, we assume $\log \rho - \gamma g_c + g_i + \frac{1}{2}(\gamma^2 \sigma_c^2 - 2\gamma \omega_{ci} \sigma_c \sigma_i + \sigma_i^2) < 0$ so that the agent's consumption–portfolio decision is well behaved at $t = \infty$.

where for $z_{i,t} \leq 1$,

$$\hat{v}(R_{i,t+1}, z_{i,t}) = \begin{cases} R_{i,t+1} - R_{f,t} & \text{for } R_{i,t+1} \geq z_{i,t} R_{f,t} \\ (z_{i,t} R_{f,t} - R_{f,t}) + \lambda(R_{i,t+1} - z_{i,t} R_{f,t}) & \text{for } R_{i,t+1} < z_{i,t} R_{f,t} \end{cases}, \quad (23)$$

and for $z_{i,t} > 1$,

$$\hat{v}(R_{i,t+1}, z_{i,t}) = \begin{cases} R_{i,t+1} - R_{f,t} & \text{for } R_{i,t+1} \geq R_{f,t} \\ \lambda(z_{i,t})(R_{i,t+1} - R_{f,t}) & \text{for } R_{i,t+1} < R_{f,t} \end{cases}. \quad (24)$$

We prove this formally in the Appendix. At a less formal level, equation (22) follows directly from the agent's Euler equation for optimality at equilibrium, derived using standard perturbation arguments,

$$1 = \rho E_t \left[R_{i,t+1} \left(\frac{\bar{C}_{t+1}}{\bar{C}_t} \right)^{-\gamma} \right] + b_0 \rho E_t [\hat{v}(R_{i,t+1}, z_{i,t})], \quad \forall i. \quad (25)$$

The first term is the standard one that obtains in an economy where investors have power utility over consumption. However, there is now an additional term. Consuming less today and investing the proceeds in stock i exposes the investor to the risk of greater losses on that stock. Just how painful this might be is determined by the state variable $z_{i,t}$.

B. Equilibrium Prices Under Portfolio Accounting

We now compute the price $P_{i,t}$ of stock i in a second economy where investors have the preferences described in equations (9) through (15).

In the case of portfolio accounting, we need to price the portfolio of all stocks in the economy before we can price any one stock. This portfolio is a claim to the aggregate dividend, which follows the process

$$\log \left(\frac{D_{t+1}}{D_t} \right) = g_p + \sigma_p \varepsilon_{t+1}, \quad (26)$$

with $\varepsilon_{t+1} \sim N(0,1)$, i.i.d. over time, and

$$\text{corr}(\eta_t, \varepsilon_{t'}) = \begin{cases} \omega_{cp} & \text{for } t = t' \\ 0 & \text{for } t \neq t' \end{cases} \quad (27)$$

$$\text{corr}(\varepsilon_{i,t}, \varepsilon_{t'}) = \begin{cases} \omega_{ip} & \text{for } t = t' \\ 0 & \text{for } t \neq t' \end{cases}. \quad (28)$$

The dividend processes for stocks 1 through n in equation (16) will not in general “add up” to the aggregate dividend process in equation (26). Without additional structure, we cannot think of the n stocks as a complete list of all stocks in the portfolio. We therefore imagine that there are some other securities in the economy whose dividends are distributed in such a way that the total dividend does add up to the aggregate dividend in equation (26). For the purpose of choosing parameters, it is helpful to have a setup where the dividends of the n original stocks alone *do* add up, and we present this special case in Section III.

Our assumptions allow us to construct a Markov equilibrium in which the risk-free rate is constant and the portfolio-level state variable z_t determines the distribution of returns on all stocks. Specifically, we assume that stock i ’s price–dividend ratio is a function of z_t ,

$$f_{i,t} \equiv P_{i,t}/D_{i,t} = f_i(z_t), \quad (29)$$

and then look for an equilibrium satisfying this assumption. Given this one-factor assumption, the distribution of the stock return $R_{i,t+1}$ is determined by z_t and the function $f_i(\cdot)$ as follows:

$$\begin{aligned} R_{i,t+1} &= \frac{P_{i,t+1} + D_{i,t+1}}{P_{i,t}} = \frac{1 + P_{i,t+1}/D_{i,t+1}}{P_{i,t}/D_{i,t}} \frac{D_{i,t+1}}{D_{i,t}} \\ &= \frac{1 + f_i(z_{t+1})}{f_i(z_t)} \frac{D_{i,t+1}}{D_{i,t}} = \frac{1 + f_i(z_{t+1})}{f_i(z_t)} e^{g_i + \sigma_i \varepsilon_{i,t+1}}. \end{aligned} \quad (30)$$

As in the first economy, the price of stock i can change because of dividend news $\varepsilon_{i,t+1}$ or because its price–dividend ratio $f_{i,t}$ changes. The key difference is that changes in this ratio are not driven by a stock-level state variable $z_{i,t}$, but by the portfolio-level state variable z_t , which tracks prior gains and losses on the overall portfolio. Recent gains (losses) on the portfolio make the investor perceive the entire portfolio as less (more) risky, changing the price–dividend ratio of every stock in the portfolio simultaneously.

In equilibrium, and under rational expectations about stock returns and aggregate consumption levels, the agents in our economy must find it optimal to hold the market supply of zero units of the risk-free asset and one unit of each stock at all times, and to consume their labor income and the aggregate dividend stream every period. The proposition below characterizes the equilibrium.

PROPOSITION 2: *For the preferences in equations (9) through (15), necessary and sufficient conditions for a one-factor Markov equilibrium are*

$$R_f = \rho^{-1} e^{\gamma g_c - \gamma^2 \sigma_c^2 / 2}, \quad (31)$$

and

$$\begin{aligned}
 1 = & \rho e^{g_i - \gamma g_c + (1/2)\gamma^2 \sigma_c^2 (1 - \omega_{cp}^2) + (1/2)\sigma_i^2 (1 - \omega_{ip}^2) - \gamma \sigma_c \sigma_i (\omega_{ci} - \omega_{cp} \omega_{ip})} \\
 & \times E_t \left[\frac{1 + f_i(z_{t+1})}{f_i(z_t)} e^{(\sigma_i \omega_{ip} - \gamma \sigma_c \omega_{cp}) \varepsilon_{t+1}} \right] \\
 & + b_0 \rho E_t \left[\tilde{v} \left(\frac{1 + f_i(z_{t+1})}{f_i(z_t)} e^{g_i + \sigma_i \omega_{ip} \varepsilon_{t+1} + (1/2)\sigma_i^2 (1 - \omega_{ip}^2)}, R_{t+1}, z_t \right) \right],
 \end{aligned} \tag{32}$$

where for $z_t \leq 1$,

$$\begin{aligned}
 & \tilde{v}(R_{i,t+1}, R_{t+1}, z_t) \\
 & = \begin{cases} R_{i,t+1} - R_{f,t} & \text{for } R_{t+1} \geq z_t R_{f,t} \\ (z_t R_{f,t} - R_{f,t}) + \lambda(R_{i,t+1} - z_t R_{f,t}) & \text{for } R_{t+1} < z_t R_{f,t} \end{cases},
 \end{aligned} \tag{33}$$

and for $z_t > 1$,

$$\tilde{v}(R_{i,t+1}, R_{t+1}, z_t) = \begin{cases} R_{i,t+1} - R_{f,t} & \text{for } R_{t+1} \geq R_{f,t} \\ \lambda(z_t)(R_{i,t+1} - R_{f,t}) & \text{for } R_{t+1} < R_{f,t} \end{cases}. \tag{34}$$

The gross return R_{t+1} on the overall portfolio is given by

$$R_{t+1} = \frac{1 + f(z_{t+1})}{f(z_t)} e^{g_p + \sigma_p \varepsilon_{t+1}}, \tag{35}$$

where $f(\cdot)$ satisfies

$$\begin{aligned}
 1 = & \rho e^{g_p - \gamma g_c + (1/2)\gamma^2 \sigma_c^2 (1 - \omega_{cp}^2)} E_t \left[\frac{1 + f(z_{t+1})}{f(z_t)} e^{(\sigma_p - \gamma \omega_{cp} \sigma_c) \varepsilon_{t+1}} \right] \\
 & + b_0 \rho E_t \left[\hat{v} \left(\frac{1 + f(z_{t+1})}{f(z_t)} e^{g_p + \sigma_p \varepsilon_{t+1}}, z_t \right) \right],
 \end{aligned} \tag{36}$$

and where $\hat{v}$ is defined in Proposition 1. Moreover, $f(\cdot)$ and $f_i(\cdot)$ must be self-consistent:

$$f(z_t) = \sum_i \frac{D_{i,t}}{D_t} f_i(z_t). \tag{37}$$

We prove this formally in the Appendix. At a less formal level, equation (32) follows from the agent's Euler equation for optimality,

$$1 = \rho E_t \left[R_{i,t+1} \left(\frac{\bar{C}_{t+1}}{\bar{C}_t} \right)^{-\gamma} \right] + b_0 \rho E_t [\tilde{v}(R_{i,t+1}, R_{t+1}, z_t)], \quad \forall i. \quad (38)$$

The first term is standard. The second term reflects the fact that consuming less today and investing the proceeds in stock i exposes the investor to the risk of greater losses on that stock. How painful this is is now determined by the portfolio-level state variable z_t .

Note that there is now an extra step in computing the price–dividend ratio of an individual stock. Under portfolio accounting, the price behavior of individual stocks depends on the behavior of the overall portfolio. This means that we first need to calculate the equilibrium portfolio return R_{t+1} in equations (35) and (36), and then use this in equation (32) governing the individual stock return.

III. Numerical Results and Intuition

In this section, we solve equations (22) and (32) for the price–dividend ratio of an individual stock $f_i(\cdot)$ under individual stock accounting and portfolio accounting, respectively, and then use simulated data to study the properties of equilibrium stock returns in each case.

A. Parameter Values

Table I summarizes our choice of parameters. We divide the table into two panels to separate the two types of parameters: those that determine the distributions of consumption and dividend growth, and those that determine investor preferences.

For the mean g_c and standard deviation σ_c of log consumption growth, we follow Cecchetti, Lam, and Mark (1990), who obtain $g_c = 1.84$ percent and $\sigma_c = 3.79$ percent from a time series of annual data from 1889 to 1985.

In principle, specifying parameters for individual stock and aggregate dividend growth is a daunting task. Equations (16) and (26) show that we need g_i , σ_i , ω_{ci} , and ω_{ip} for each stock, ω_{ij} for each pair of stocks, and g_p , σ_p , and ω_{cp} for the aggregate portfolio, a total of $(n^2/2) + (7n/2) + 3$ parameters. Fortunately, it turns out that with two simplifying assumptions, we can specify the dividend processes with just four parameters, and yet still convey most of the important economics. Essentially, we take dividend growth to be identically distributed across all stocks, a restriction that we relax in Section III.E.

Our first assumption is that the mean and standard deviation of log dividend growth is the same for all stocks:

$$g_i = g, \quad \sigma_i = \sigma, \quad \forall i. \quad (39)$$

Table I
Parameter Values

Panel A lists the parameters that determine the distributions of consumption and dividend growth. Specifically, g_c and σ_c are the mean and standard deviation of log consumption growth; g_i , σ_i , ω_{ci} , and ω_{ij} are the mean, standard deviation, correlation with log consumption growth, and correlation with stock j 's log dividend growth of stock i 's log dividend growth, respectively; and g_p , σ_p , ω_{cp} , and ω_{ip} are the mean, standard deviation, correlation with log consumption growth, and correlation with stock i 's log dividend growth of log aggregate dividend growth. Panel B lists the parameters that determine investor preferences: γ governs the curvature of utility over consumption, ρ is the time discount rate, λ determines how keenly losses are felt relative to gains when there are no prior gains and losses, k determines how much more painful losses are when they come on the heels of other losses, b_0 determines the relative importance of loss aversion in the investor's preferences, and η determines how long prior gains and losses affect the investor.

Panel A: Consumption and Dividend Parameters		
g_c	1.84%	
σ_c	3.79%	
g_i	-0.91%,	$\forall i$
σ_i	25.0%,	$\forall i$
ω_{ci}	0.072,	$\forall i$
ω_{ij}	0.23,	$\forall i, j$
g_p	1.5%	
σ_p	12.0%	
ω_{cp}	0.15	
ω_{ip}	0.48,	$\forall i$
Panel B: Preference Parameters		
γ	1.0	
ρ	0.98	
λ	2.25	
k	3.0	
b_0	0.45	
η	0.9	

Second, we assume a simple factor structure for individual stock dividend growth innovations:

$$\varepsilon_{i,t} = \omega_p \varepsilon_t + \hat{\varepsilon}_{i,t} \sqrt{1 - \omega_p^2}. \quad (40)$$

In words, the cash-flow shock to stock i has one component due to the aggregate dividend innovation ε_t introduced in equation (26), and one idiosyncratic component, $\hat{\varepsilon}_{i,t} \sim \text{i.i.d. } N(0,1)$. The relative importance of the two components is controlled by a new parameter, ω_p . The idiosyncratic component is orthogonal to consumption growth shocks, aggregate dividend growth shocks, and the idiosyncratic shocks on other stocks:

$$\text{corr}(\hat{\varepsilon}_{i,t}, \eta_t) = \text{corr}(\hat{\varepsilon}_{i,t}, \varepsilon_t) = \text{corr}(\hat{\varepsilon}_{i,t}, \hat{\varepsilon}_{j,t}) = 0, \quad \forall i, j. \quad (41)$$

This immediately implies

$$\omega_{ci} = \omega_p \omega_{ep}, \quad \forall i, \quad (42)$$

$$\omega_{ij} = \omega_p^2, \quad \forall i, j, \quad (43)$$

$$\omega_{ip} = \omega_p, \quad \forall i. \quad (44)$$

Another attractive feature of this simple factor structure is that in the limit, as we add more and more stocks, the growth of their total dividend is also i.i.d lognormal:¹³

$$\lim_{n \rightarrow \infty} \frac{\sum_{i=1}^n D_{i,t+1}}{\sum_{i=1}^n D_{i,t}} \rightarrow e^{g + (1/2)\sigma^2(1 - \omega_p^2) + \sigma\omega_p \varepsilon_{t+1}}. \quad (45)$$

This means that we *can* think of the n stocks as being an exhaustive list of all securities, with their total dividend equaling the aggregate dividend in (26), $D_t = \sum_{i=1}^n D_{i,t}$. Comparing equation (45) with equation (26), we obtain¹⁴

$$g_p = g + \frac{1}{2} \sigma^2 (1 - \omega_p^2), \quad (46)$$

$$\sigma_p = \sigma \omega_p. \quad (47)$$

Equations (39), (42) through (44), and (46) through (47) show that the entire structure of dividend growth can be determined from g_p , σ_p , σ , and ω_{ep} alone. We choose these four quantities as the basis variables rather than any other four because they can be estimated in a relatively straightforward manner. First, we estimate the mean and standard deviation of aggregate dividend growth using NYSE data from 1925 to 1995 from CRSP, which

¹³ For an economy with a finite horizon from $t = -T$ to $t = T$, this limiting behavior is based on an argument similar to the law of large numbers. Anderson (1991) and Green (1989) provide technical details. Our stationary economy can then be thought of as the limit as the time horizon goes to infinity.

¹⁴ The assumptions in this section also allow us to satisfy the self-consistency condition in Proposition 2. Under our assumptions, if $f(\cdot)$ solves equation (36), then it also solves equation (32) for all i . Therefore $f_i(\cdot) = f(\cdot)$, $\forall i$, and so the self-consistency condition is satisfied.

gives $g_p = 0.015$ and $\sigma_p = 0.12$. The correlation between shocks to consumption growth and dividend growth, ω_{cp} , we take from Campbell (2000), who estimates it in the neighborhood of 0.15.

We set the volatility σ of individual stock dividend growth to 0.25. A direct calculation of the value-weighted average volatility of real dividend growth for firms in the COMPUSTAT database suggests that this is a reasonable benchmark level.¹⁵ Further confirmation comes from Vuolteenaho (1999), who estimates the cash-flow news volatility of an individual stock, equal weighted across stocks, to be 32 percent. Since smaller firms have more volatile cash flows, 25 percent may be a better estimate of value-weighted cash-flow volatility.

Panel A in Table I shows what these values imply for the remaining parameters governing the dividend processes.¹⁶ The preference parameters are summarized in Panel B of the table. We choose the curvature γ of utility over consumption and the time discount factor ρ so as to produce a sensibly low value for the risk-free rate. Given the values of g_c and σ_c , equation (21) shows that $\gamma = 1.0$ and $\rho = 0.98$ set the risk-free interest rate to $R_f - 1 = 3.86$ percent.

The value of λ determines how keenly losses are felt relative to gains in the case where the investor has no prior gains or losses. We take $\lambda = 2.25$, the value Tversky and Kahneman (1992) estimate by offering subjects isolated gambles in experimental settings.

The parameter k determines how much more painful losses are when they come on the heels of other losses. We choose $k = 3$. To interpret this, suppose that the state variable $z_{i,t}$ is initially equal to one, and that stock i then experiences a sharp fall of 10 percent. From equation (8) with $\eta \approx 1$, this means that $z_{i,t}$ goes up by about 0.1, to 1.1. From equation (7), any additional losses will now be penalized at $2.25 + 3(0.1) = 2.55$, a slightly more severe penalty.

The parameter b_0 determines the relative importance of the loss aversion term in the investor's preferences. We set $b_0 = 0.45$. One way to think about b_0 is to compare the disutility of losing a dollar on a stock investment with the disutility of having to consume a dollar less. When computed at equilibrium, the ratio of these two quantities equals $b_0\rho\lambda$. Our parameter choices therefore make the psychological disutility of losing the \$1 roughly equal in magnitude to the consumption disutility.

Finally, η arises in the definition of the state variable dynamics. It controls the persistence of the state variable $z_{i,t}$, which in turn controls the autocorrelation of price-dividend ratios. We find that $\eta = 0.9$ brings this autocorrelation close to its empirical value.

¹⁵ More precisely, we take all stocks in the annual COMPUSTAT database for which at least 11 consecutive years of dividend data are recorded, compute real dividend growth volatility for each, and then calculate the average volatility, weighted by firm size.

¹⁶ Note that while the mean *log* dividend growth g_i is negative, mean *simple* dividend growth equals $\exp[-0.0091 + (0.25^2/2)] - 1 = 2.24$ percent, a positive number.

B. Methodology

For the case of individual stock accounting, we use an iterative technique to solve equation (22) for the price–dividend ratio $f_i(\cdot)$ of an individual stock. The only difficulty is that the state variable $z_{i,t}$ is endogenous: It tracks prior gains and losses which depend on past returns, themselves endogenous. To deal with this, we use the following procedure. We guess a solution to equation (22), $f_i^{(0)}$ say, and then construct a function $z_{i,t+1} = h_i^{(0)}(z_{i,t}, \varepsilon_{i,t+1})$ that solves

$$R_{i,t+1} = \frac{1 + f_i(z_{i,t+1})}{f_i(z_{i,t})} e^{g_i + \sigma_i \varepsilon_{i,t+1}} \quad (48)$$

and

$$z_{i,t+1} = \eta \left(z_{i,t} \frac{\bar{R}}{R_{i,t+1}} \right) + (1 - \eta)(1) \quad (49)$$

simultaneously for this particular $f_i = f_i^{(0)}$. Given $h_i^{(0)}$, we get a new candidate solution $f_i^{(1)}$ through the recursion

$$\begin{aligned} 1 = & \rho e^{g_i - \gamma g_c + (1/2)\gamma^2 \sigma_c^2 (1 - \omega_{ci}^2)} E_t \left[\frac{1 + f_i^{(j)}(z_{i,t+1})}{f_i^{(j+1)}(z_{i,t})} e^{(\sigma_i - \gamma \omega_{ci} \sigma_c) \varepsilon_{i,t+1}} \right] \\ & + b_0 \rho E_t \left[\hat{v} \left(\frac{1 + f_i^{(j)}(z_{i,t+1})}{f_i^{(j+1)}(z_{i,t})} e^{g_i + \sigma_i \varepsilon_{i,t+1}}, z_{i,t} \right) \right]. \end{aligned} \quad (50)$$

With $f_i^{(1)}$ in hand, we calculate a new $h_i = h_i^{(1)}$ that solves equations (48) and (49) for $f_i = f_i^{(1)}$. This $h_i^{(1)}$ gives us a new candidate $f_i = f_i^{(2)}$ from equation (50). We continue this process until we obtain convergence, $f_i^{(j)} \rightarrow f_i$, $h_i^{(j)} \rightarrow h_i$. Figure 2 shows the price–dividend ratio $f_i(\cdot)$ that corresponds to the parameter values in Table I. Its precise shape will be explained in more detail later.

With the price–dividend ratio $f_i(\cdot)$ in hand, we use simulated data to see how returns behave in equilibrium. We simulate dividend shocks $\{\varepsilon_{i,t}\}$ for $n = 500$ stocks and for 10,000 time periods subject to the specification in equation (16) and the parameters in Table I. We then apply the price–dividend function $f_i(\cdot)$ to this dividend data to see what realized returns look like. More precisely, we use the $z_{i,t+1} = h_i(z_{i,t}, \varepsilon_{i,t+1})$ function described above to generate the series of $z_{i,t}$ implied by the dividend shocks and then set the return of stock i between time t and $t + 1$ equal to

$$R_{i,t+1} = \frac{1 + f_i(z_{i,t+1})}{f_i(z_{i,t})} e^{g_i + \sigma_i \varepsilon_{i,t+1}}. \quad (51)$$




Figure 2. Price-dividend ratio of stock i under individual stock accounting and portfolio accounting. The two price-dividend curves are presented on the same graph for ease of comparison, but they are functions of different variables. Under individual stock accounting, the price-dividend ratio is a function of $z_{i,t}$, which measures prior gains and losses on stock i . Under portfolio accounting, it is a function of z_t , which measures prior gains and losses on the investor's overall portfolio of stocks.

This gives n time series of individual stock returns. We can then compute moments—standard deviation, say—for each stock, and then average these moments across different stocks. This provides a sense of how the “typical” stock behaves, and we report the results of such calculations later in Section III.

We can also use our n time series of individual stock returns to compute an equal-weighted average

$$R_{p,t+1} = \frac{1}{n} \sum_{i=1}^n R_{i,t+1}, \quad (52)$$

which can be interpreted as the aggregate stock return.¹⁷

¹⁷ The aggregate return $R_{p,t+1}$ differs from the aggregate return R_{t+1} described earlier only in that it is equal-weighted rather than value-weighted.

For the case of portfolio accounting, we start out by using iteration in equation (36) to compute the aggregate price–dividend ratio. As before, we iterate between guesses $f = f^{(j)}$ and functions $z_{t+1} = h^{(j)}(z_t, \varepsilon_{t+1})$ that solve

$$R_{t+1} = \frac{1 + f(z_{t+1})}{f(z_t)} e^{g_p + \sigma_p \varepsilon_{t+1}} \quad (53)$$

and

$$z_{t+1} = \eta \left(z_t \frac{\bar{R}}{R_{t+1}} \right) + (1 - \eta)(1) \quad (54)$$

simultaneously for $f = f^{(j)}$. Once this process converges with $f^{(j)} \rightarrow f$ and $h^{(j)} \rightarrow h$, we take the resulting functions $f(\cdot)$ and $z_{t+1} = h(z_t, \varepsilon_{t+1})$ and iterate in equation (32) over guesses $f_i^{(j)}(\cdot)$ for stock i 's price–dividend ratio, converging eventually to the solution $f_i(\cdot)$.¹⁸ Figure 2 shows the price–dividend ratio $f_i(\cdot)$ that we obtain for the parameter values in Table I. We display it on the same graph as the price–dividend ratio obtained earlier in the individual stock accounting case, but it is important to note that the two curves are plotted against different state variables: against $z_{i,t}$ for individual stock accounting, and against z_t for portfolio accounting.

Simulation then illustrates the behavior of individual stock returns. We again generate dividend shocks $\{\varepsilon_{i,t}\}$ for $n = 500$ stocks and for the portfolio, $\{\varepsilon_t\}$, over 10,000 time periods. The function $z_{t+1} = h(z_t, \varepsilon_{t+1})$ generates the time series for the aggregate state variable z_t implied by these $\{\varepsilon_t\}$. The time series of returns for stock i is then given by

$$R_{i,t+1} = \frac{1 + f_i(z_{t+1})}{f_i(z_t)} e^{g_i + \sigma_i \varepsilon_{i,t+1}}, \quad (55)$$

while the aggregate return is measured by

$$R_{p,t+1} = \frac{1}{n} \sum_{i=1}^n R_{i,t+1}. \quad (56)$$

C. Equilibrium Returns Under Individual Stock Accounting

Table II summarizes the properties of individual and aggregate stock returns in simulated data from three economies: one in which investors use individual stock accounting, another in which they use portfolio accounting, and for comparison, a third economy in which investors experience no loss

¹⁸ The way $f(\cdot)$ enters equation (32) is through the portfolio return R_{t+1} in equations (33) and (34). Note that R_{t+1} depends on $f(\cdot)$ as shown in equation (35).

Table II
Properties of Asset Returns in Simulated Data

We report the properties of asset returns in simulated data from three economies. In the first economy, labeled “Individual Stock Accounting,” investors are loss averse over individual stock fluctuations; in the second, labeled “Portfolio Accounting,” they are loss averse over portfolio fluctuations; and in the third, labeled “No Loss Aversion,” they have power utility over consumption levels. Panel A reports the properties of the typical individual stock return: its mean in excess of the risk-free rate, its standard deviation, the R^2 in a regression of four-year cumulative log returns on the lagged dividend yield, and the average contemporaneous return correlation with other stocks. Panel B reports the risk-free rate and the properties of the aggregate stock return: its mean in excess of the risk-free rate, its standard deviation, the R^2 in a regression of four-year cumulative log returns on the lagged dividend yield, and its correlation with consumption growth. Panel C reports the cross-sectional features of the data. Each year, we sort stocks into deciles based on their price–dividend ratio, and measure the returns of the top and bottom decile portfolios over the next year. The value premium is the time series mean of the difference in returns between the two deciles. Every three years, we sort stocks into deciles based on their three-year prior return, and measure the average annual returns of the top and bottom deciles over the next three years. The De Bondt–Thaler premium is the time series mean of the difference in average returns between the two deciles over all nonoverlapping periods.

	Individual Stock Accounting	Portfolio Accounting	No Loss Aversion
Panel A: Properties of Individual Stock Returns			
Excess mean $E(R) - R_f$	6.7%	2.2%	0.1%
Standard deviation σ_R	40.6%	29.4%	26.5%
Predictability R^2	4.1%	0.4%	0.0%
Correlation ω_R	0.23	0.34	0.23
Panel B: Properties of Aggregate Asset Returns			
Aggregate stock return			
Excess mean $E(R_M) - R_f$	6.7%	2.2%	0.1%
Standard deviation σ_M	19.0%	17.2%	12.7%
Predictability $R^2(M)$	5.9%	2.1%	0.0%
Correlation with consumption growth	0.15	0.15	0.15
Risk-free rate R_f	3.86%	3.86%	3.86%
Panel C: Properties of the Cross Section			
Value premium	17.9%	—	—
De Bondt–Thaler premium	11.1%	0.0%	0.0%

aversion at all. Specifically, in this third economy, investors have the preferences in equation (2) but with $b_0 = 0$; in other words, they have power utility over consumption with $\gamma = 1$ and $\rho = 0.98$.

Panel A in Table II reports time series properties of *individual stock* returns; Panel B describes the time series properties of *aggregate* returns; finally, Panel C summarizes the *cross-sectional* patterns in average returns. As described in Section III.B, the time series results for individual stocks come from computing the relevant moment for each stock in the simulated

sample and then averaging across stocks. All values are expressed in annual terms. In this section, as well as in Sections III.D and III.E, we lay out the results and explain the intuition behind them. Section IV discusses the broader implications of our findings.

As we present the results, it may be helpful to keep in mind the main empirical findings that have been documented. In the time series, the returns of a typical individual stock have a high mean, are excessively volatile, and are slightly predictable using lagged variables. The time series of aggregate stock returns displays the same properties. In the cross section, there is a substantial value premium, in that stocks with low ratios of price to fundamentals have higher average returns.¹⁹ The well-known difficulties that the “No Loss Aversion” model faces in explaining these facts are clearly illustrated in Table II: The equity premium is tiny, there is no excess volatility to speak of, no time series predictability, and no value premium in the cross section.

In our first set of results, we study an economy in which investors do individual stock accounting. We look first at the time series implications. The typical individual stock has a high average excess return, $E(R) - R_f = 6.7$ percent. It is volatile, and in particular, excessively volatile: Its standard deviation, $\sigma_R = 40.6$ percent, is higher than the standard deviation of underlying dividend growth, $\sigma = 25$ percent. Its returns are also slightly predictable: A regression of four-year cumulative log returns on the lagged dividend yield,

$$\log(R_{i,t+1}R_{i,t+2}R_{i,t+3}R_{i,t+4}) = \alpha + \beta \frac{D_{i,t}}{P_{i,t}} + v_{i,t+4},$$

produces an R^2 of 4.1 percent.

The last variable in Panel A, ω_R , measures the average contemporaneous correlation of individual stocks; we calculate it to be 0.23 in a way that we explain shortly.

The next panel shows that the aggregate stock market also has a high excess average return, $E(R_M) - R_f = 6.7$ percent, and is excessively volatile, with a standard deviation of $\sigma_M = 19$ percent that exceeds the 12 percent standard deviation of aggregate dividend growth. Aggregate stock returns are also slightly predictable: A regression of four-year cumulative log aggregate returns on the lagged aggregate dividend yield gives an R^2 of 5.9 percent.²⁰

¹⁹ LeRoy and Porter (1981) and Shiller (1981) find excess volatility in aggregate stock returns, while Campbell and Shiller (1988) and Fama and French (1988) show that the dividend yield has predictive power for future aggregate returns. Vuolteenaho (1999) documents the excess volatility and time series predictability of firm level stock returns. The value premium was originally noted by Basu (1983) and Rosenberg, Reid, and Lanstein (1985); Fama and French (1992) provide more recent evidence.

²⁰ Note that Table II uses the notation $R^2(M)$ to distinguish this from the R^2 obtained earlier in the individual stock return regression.

The market standard deviation σ_M helps us measure the average correlation ω_R between stocks that we reported earlier: We compute it as $(\sigma_M/\sigma_R)^2$. This calculation is exact in the limit as the number of stocks $n \rightarrow \infty$, if all stocks have the same standard deviation and correlation with one another, as they do in our simple economy. This follows because

$$\sigma_M^2 = \lim_{n \rightarrow \infty} \text{Var} \left(\frac{R_{1,t+1} + \dots + R_{n,t+1}}{n} \right) = \lim_{n \rightarrow \infty} \left(\frac{\sigma_R^2}{n} + \left(1 - \frac{1}{n} \right) \sigma_R^2 \omega_R \right) = \sigma_R^2 \omega_R.$$

Panel C in Table II describes the cross-sectional features of individual stock returns. Our simulated data produces a value premium: A scaled-price variable—in our case, the price–dividend ratio—has predictive power for the cross section of average returns. Each year, we sort stocks into deciles based on this ratio, and measure the returns of the top and bottom decile portfolios over the next year. The time series mean of the difference in returns between the two deciles is a very substantial 17.9 percent.

Our data also replicates a well-known study of De Bondt and Thaler (1985), which finds that long-term prior losing stocks on average outperform long-term prior winning stocks. Every three years, we sort stocks into deciles based on their three-year prior return, and measure the average annual returns of the top and bottom deciles over the next three years. The time series mean of the difference in average returns between the two deciles over all nonoverlapping periods in our simulated data is 11.1 percent.

Many of the effects we obtain under individual stock accounting derive from a single source, namely a discount rate for individual stocks that changes as a function of the stock's past performance. If a stock has had good recent performance, the investor gets utility from this gain, and becomes less concerned about future losses on the stock because any losses will be cushioned by the prior gains. In effect, the investor perceives the stock to be less risky than before and discounts its future cash flows at a lower rate. Conversely, if one of his stocks performs dismally, he finds this painful and becomes more sensitive to the possibility of further losses on the stock. In effect, he views the stock as riskier than before and raises its discount rate.

This changing discount rate has many implications. It gives individual stocks some time series predictability: A lower discount rate pushes up the price–dividend ratio and leads to lower subsequent returns, which means that the lagged price–dividend ratio can predict returns. It makes stock returns more volatile than underlying cash flows: A high cash flow pushes the stock price up, but this prior gain also lowers the discount rate on the stock, pushing the stock price still higher. It also generates a value premium in the cross section: In our economy, a stock with a high price–dividend ratio (a growth stock) is often one that has done well in the past, accumulating prior gains for the investor who then views it as less risky and requires a lower average return. A stock with a low price–dividend ratio (a value stock) has often had dismal prior performance, burning the investor, who now views it as riskier, and requires a higher average return.

The high equity premia we obtain under individual stock accounting derive from a different source: Since the investor is loss averse over individual stock fluctuations, he dislikes the frequent losses that individual stocks often produce, and charges a high average return as compensation. Other papers, such as Benartzi and Thaler (1995) and Barberis et al. (2001) have also suggested loss aversion as a way of understanding a high equity premium. The effect we obtain here, though, is one level stronger than in those earlier papers, since the investor is now loss averse over individual stock fluctuations rather than over the less dramatic fluctuations in the diversified aggregate market.

This intuition also explains why the price–dividend function in Figure 2 is downward sloping. A lower value of $z_{i,t}$ means that the investor has accumulated prior gains on stock i . Since he is now less concerned about future losses on this stock, he perceives it to be less risky and is therefore willing to pay a higher price for it per unit of cash flow.

D. Equilibrium Returns Under Portfolio Accounting

Our next set of results shows that the investor's system of mental accounting matters a great deal for the behavior of asset prices. As we broaden the investor's frame from individual stock accounting to portfolio accounting, individual stock returns exhibit quite different features in equilibrium.

Table II shows that under portfolio accounting, the average excess return on a typical individual stock is 2.2 percent—not insubstantial, but rather lower than under individual stock accounting. At 29.4 percent, individual stock volatility is also lower than under individual stock accounting. In particular, excess volatility of returns over dividend growth is much smaller. Individual stock returns are predictable in the time series, but only slightly.

The average excess return on the aggregate market is 2.2 percent. Interestingly, aggregate returns are roughly as volatile here as they are under individual stock accounting. Since individual stocks are much less volatile here than under individual stock accounting, this must mean that stocks are more highly correlated than before, and indeed, $\omega_R = 0.34$. Finally, aggregate stock returns are slightly predictable.

We now turn to the cross section. One disadvantage of our assumption that stock-specific parameters—dividend growth mean g_i , standard deviation σ_i , and correlations with the overall portfolio ω_{ip} and with consumption ω_{ci} —are the same for all stocks is that there is no cross-sectional dispersion in price–dividend ratios in the case of portfolio accounting. This assumption will be relaxed in Section III.E. For now, the lack of dispersion means that we cannot check for a value premium in the simulated data.

We can, however, still look to see if there is a De Bondt–Thaler premium to prior losers. As Table II shows, this effect is no longer present under portfolio accounting.

The reason the results are different under portfolio accounting is that in this case, changes in discount rates on stocks are driven by fluctuations in the value of the overall portfolio. When the portfolio does well, the investor is less concerned

about losses on *any* of the stocks that he holds, since the prior portfolio gain will cushion any such losses. Effectively, he views *all* stocks as less risky. Discount rates on all stocks therefore go down simultaneously. Conversely, discount rates on all stocks go up after a prior portfolio loss.

This discount rate behavior is the key to many of the portfolio accounting results. Stock returns are less volatile here than under individual stock accounting. In the latter case, stocks are highly volatile because good cash-flow news is always accompanied by a lower discount rate, pushing the price up even more. Under portfolio accounting, good cash-flow news on a particular stock will only coincide with a lower discount rate on the stock if the portfolio as a whole does well. There is no guarantee of this, and so volatility is not amplified by as much. Since shocks to discount rates are perfectly correlated across stocks, individual stock returns are highly correlated with one another. Moreover, the De Bondt–Thaler premium disappears because a stock's past performance no longer affects its discount rate, which is now determined at the portfolio level.

Finally, while there is a substantial equity premium, it is not as large as under individual stock accounting. The investor is loss averse over portfolio level fluctuations, which are sizable but not as severe as the swings on individual stocks. The compensation for risk is therefore more moderate.

This intuition also clarifies why the price–dividend function in Figure 2 is downward sloping. A lower value of z_t means that the investor has accumulated prior gains on his portfolio. Since he is now less concerned about future losses on stock i —or indeed, on any stock—he perceives stock i to be less risky and is therefore willing to pay a higher price for it per unit of cash flow. Since he is loss averse only over portfolio fluctuations and not over individual stock fluctuations, he on average perceives stocks to be less risky, which is why the overall level of the price–dividend function is higher here than under individual stock accounting.

E. Further Cross-sectional Results

Our analysis so far has assumed that mean log dividend growth rates g_i and log dividend growth volatilities σ_i are equal across stocks. In particular, Table I shows that we have assumed $g_i = -0.0091$ and $\sigma_i = 0.25$, $\forall i$. Note that for the mean *simple* dividend growth rate G_i , this implies

$$G_i \equiv E\left(\frac{D_{i,t+1}}{D_{i,t}}\right) - 1 = e^{g_i + \sigma_i^2/2} - 1 = 0.0224, \quad \forall i$$

$$\ln(1 + G_i) = 0.0222, \quad \forall i.$$

We now relax these restrictive parameter choices. This should allow for a more realistic comparison of the cross-sectional features of the actual data with those of the simulated data.

Table III

The Cross Section of Average Returns in Simulated Data

We report the cross-sectional properties of average returns in simulated data which allows for dispersion in both the mean and the standard deviation of firm-level dividend growth rates. We consider three economies. In the first economy, labeled “Individual Stock Accounting,” investors are loss averse over individual stock fluctuations; in the second, labeled “Portfolio Accounting,” they are loss averse over portfolio fluctuations; in the third, labeled “No Loss Aversion,” they have power utility over consumption levels. Each year, we sort stocks into deciles based on their price–dividend ratio, and measure the returns of the top and bottom decile portfolios over the next year. The value premium is the time series mean of the difference in returns between the two deciles. Every three years, we sort stocks into deciles based on their three-year prior return, and measure the average annual returns of the top and bottom deciles over the next three years. The De Bondt–Thaler premium is the time series mean of the difference in average returns between the two deciles over all nonoverlapping periods.

	Individual Stock Accounting	Portfolio Accounting	No Loss Aversion
Value premium	12.6%	0.0%	0.0%
De Bondt–Thaler premium	8.2%	0.0%	0.0%

As before, we simulate dividend data for 500 different stocks over 10,000 time periods. This time, however, we draw G_i and σ_i for each stock from the following distributions, independently across different stocks:

$$\ln(1 + G_i) \sim N(0.0222, \sigma_g^2)$$

$$\sigma_i \sim N(0.25, \sigma_S^2),$$

where $\sigma_g = 0.01$ and $\sigma_S = 0.05$.²¹ In other words, we allow for dispersion in the mean and volatility of dividend growth across different stocks. We use the dispersion in dividend growth volatilities estimated from COMPUSTAT data as a guide to choosing σ_S ; σ_g is harder to estimate directly, but we find that given our other parameter values, a σ_g of 0.01 leads to a realistic cross-sectional dispersion in price–dividend ratios.

Table III repeats the cross-sectional calculations shown in the bottom panel of Table II for our new simulated data.²² We look first at the results for individual stock accounting.

Under individual stock accounting, the value premium and the De Bondt–Thaler premium to prior losers remain as strong as in Table II. The bulk of the value premium here is due to changing discount rates on individual

²¹ We implement any particular draw $(\hat{G}_i, \hat{\sigma}_i)$ by setting $\sigma_i = \hat{\sigma}_i$ and $g_i = -0.0091 + \ln(1 + \hat{G}_i) - 0.0222 + \frac{1}{2}(0.25^2 - \hat{\sigma}_i^2)$, which immediately implies $G_i = \hat{G}_i$.

²² To avoid computing the price–dividend function $f_i(\cdot)$ and transition function $h_i(\cdot)$ for all 500 draws of G_i and σ_i , we compute them for a few (G_i, σ_i) pairs and then use interpolation to approximate f_i and h_i for other values of G_i and σ_i .

stocks, exactly as described earlier. Now that we have allowed for cross-sectional dispersion in the volatility of dividend growth, there is an additional mechanism contributing to the value premium. A firm with more volatile cash-flow growth has more volatile returns, which scares the loss-averse investor into charging a higher risk premium. Since the investor is applying a high discount rate to future cash flows, the firm will have a low price-dividend ratio, thus generating a cross-sectional link between price-dividend ratios and average returns.²³

The value premium in Table III is lower than the premium in Table II. The reason for this is the following. In the Table II simulations, a stock only has a low price-dividend ratio if investors decide to assign it a high discount rate. Portfolios of low price-dividend stocks therefore earn high returns on average. In the more realistic simulations we do in Table III, a low price-dividend ratio can be a sign of a high discount rate, but it can also mean that the stock has a low g_i , in other words, that investors expect low cash-flow growth on the stock. Portfolios of low price-dividend stocks still have high returns, but the effect is diluted since the portfolio now includes stocks with low expected growth rates and only average discount rates.

Table III also presents results for the portfolio accounting case.²⁴ There is now dispersion in price-dividend ratios, which means that we can examine whether price-dividend ratios have any predictive power in the cross section. They do not; in other words, there is no value premium. The attempt to replicate De Bondt and Thaler's (1985) findings is also a failure, as it was in Table II.

Given the ability of individual stock accounting to generate a value premium, it is natural to ask whether we are also able to replicate the finding of Fama and French (1993), namely that the value premium can be captured by a multifactor model that includes as a factor a portfolio formed by ranking stocks on their scaled-price ratios—the so-called HML portfolio.

To study this, we simplify the cash-flow structure by once again equating the mean and volatility of dividend growth across all stocks, but also enrich it by allowing firm level cash flows to be driven not just by one factor, but by many factors. These additional factors can be thought of as industry shocks. Suppose that each of the n stocks falls into one of a small number of industry sectors, labelled $s = 1, \dots, S$. Suppose that the cash-flow shock to stock i now has the structure

$$\varepsilon_{i,t+1} = \omega_p \varepsilon_{t+1} + \omega_s \varepsilon_{s,t+1} + \hat{\varepsilon}_{i,t+1} \sqrt{1 - \omega_p^2 - \omega_s^2}. \quad (57)$$

²³ Our claim that the bulk of the value premium is due to changing discount rates is based on a comparison of the results in Table III to those for an economy in which there is dispersion *only* in the mean dividend growth rate across firms and not in the volatility of dividend growth. These results are not shown here but are available from the authors on request.

²⁴ Now that we have allowed for dispersion in g_i and σ_i , we cannot be certain that the self-consistency condition in Proposition 2 is satisfied. The portfolio accounting results in Table III should therefore be viewed as approximate.

In addition to the market-wide shock ε_{t+1} and the idiosyncratic shock $\hat{\varepsilon}_{i,t+1}$, there is sector-level cash-flow shock $\varepsilon_{s,t+1}$, where s is the sector to which stock i belongs. All shocks are distributed $N(0,1)$, i.i.d. over time, and their relative importance is determined by ω_p and ω_s . If we assume for simplicity that these shocks are orthogonal to one another, the correlation ω_{ij} between the cash-flow shocks on stock i and j will satisfy

$$\omega_{ij} = \omega_p^2, \quad \text{for } i \neq j \text{ and } i, j \text{ in different sectors,}$$

$$\omega_{ij} = \omega_p^2 + \omega_s^2, \quad \text{for } i \neq j \text{ and } i, j \text{ in the same sector,}$$

$$\omega_{ij} = 1, \quad \text{for } i = j.$$

We simulate dividend data on 500 stocks over 10,000 time periods using this correlation structure, and then compute the resulting stock prices in an economy where investors do individual stock accounting. We take $\omega_p = 0.48$ as before, and $\omega_s = 0.5$.

To investigate Fama and French's (1993) finding, we split our simulated data into two equal subsamples. In the first subsample, we create the portfolios whose average returns we want to explain. Each year, we sort stocks into quintiles based on their price-dividend ratio and record the equal-weighted return $R_{j,t}^{P/D}$, $j = 1, \dots, 5$, of each quintile over the next year; in particular, $R_{1,t}^{P/D}$ is the return on the first quintile, containing stocks with the lowest price-dividend ratios. Repeating this each period produces long time series of returns on the five portfolios. Table IV reports each portfolio's average return over time, $\bar{R}_j^{P/D}$. Of course, the dispersion in the average returns across the five portfolios is what we refer to as the value premium.

Table IV shows the results of time series regressions of the excess returns on these five portfolios on the excess market return,

$$R_{j,t}^{P/D} - R_f = \alpha_j + \beta_{j,1}(R_{p,t} - R_f) + u_{j,t}. \quad (58)$$

The intercepts are large, replicating Fama and French's (1993) finding that the CAPM is *not* able to explain the average returns on portfolios sorted by scaled-price ratios.

In our second subsample, we construct a counterpart to Fama and French's (1993) HML factor: Each year, we compute the return on the portfolio of stocks with price-dividend ratios in the lowest quartile that period minus the return on the portfolio of stocks with price-dividend ratios in the highest quartile that period. We denote this difference F_t .²⁵

²⁵ The reason we split the data into two subsamples is so that $R_{j,t}^{P/D}$ and F_t can be computed in different samples, thus ensuring that our regressions do not pick up any spurious correlation.

Table IV
Time Series Regressions of Returns of Portfolios Formed
on Price–Dividend Ratios on the Excess Market Return
and on HML-type Portfolio Returns

We split our simulated data into two equal subsamples. Using the first subsample, each year, we sort stocks into quintiles based on their price–dividend ratio and record the equal-weighted return $R_{j,t}^{P/D}$, $j = 1, \dots, 5$, of each quintile over the next year; $j = 1$ is the quintile containing stocks with the lowest price–dividend ratios. $\bar{R}_j^{P/D}$ is the average return of portfolio j over time. Panel A reports the result of time series regressions of excess portfolio returns on excess market returns,

$$R_{j,t}^{P/D} - R_f = \alpha_j + \beta_{j,1}(R_{p,t} - R_f) + u_{j,t}.$$

Using the second subsample, we create a portfolio similar to Fama and French’s (1993) HML factor. Each year, we compute the return on the portfolio of stocks with price–dividend ratios in the lowest quartile that period minus the return on the portfolio of stocks with price–dividend ratios in the highest quartile that period. We denote this difference F_t . Panel B reports the result of time series regressions of excess portfolio returns on excess market returns and F_t ,

$$R_{j,t}^{P/D} - R_f = \alpha_j + \beta_{j,1}(R_{p,t} - R_f) + \beta_{j,2}F_t + u_{j,t}.$$

Panel A: One-factor Model				
j	$\bar{R}_j^{P/D} - 1$	α_j	$\beta_{j,1}$	
1	19.1%	7.3%	1.18	
2	12.9%	1.7%	1.09	
3	9.4%	−1.2%	1.01	
4	6.7%	−3.3%	0.92	
5	4.7%	−4.5%	0.81	
Panel B: Multifactor Model				
j	$\bar{R}_j^{P/D} - 1$	α_j	$\beta_{j,1}$	$\beta_{j,2}$
1	19.1%	2.2%	1.03	0.47
2	12.9%	0.1%	1.03	0.15
3	9.4%	−0.1%	1.02	−0.05
4	6.7%	−0.1%	0.99	−0.22
5	4.7%	−0.3%	0.94	−0.39

Table IV shows that when we include F_t in the time series regressions,

$$R_{j,t}^{P/D} - R_f = \alpha_j + \beta_{j,1}(R_{p,t} - R_f) + \beta_{j,2}F_t + u_{j,t}, \tag{59}$$

the intercepts fall in magnitude, replicating Fama and French’s (1993) result that loadings on this additional factor can help capture average returns on the five portfolios. Portfolio 1 (portfolio 5) loads positively (negatively) on F_t because the stocks it contains share a common industry factor with stocks whose returns enter positively (negatively) into the construction of F_t . Moreover, portfolio 1 has a higher average return than portfolio 5. Loadings on F_t therefore line up with average returns.

IV. Discussion

We draw a number of conclusions from the results in Tables II, III, and IV. First, both kinds of narrow framing may shed light on certain aspects of the data. Under both individual stock accounting and portfolio accounting, firm level returns have a high mean, are excessively volatile, and are predictable in the time series using lagged variables. In both cases, aggregate stock returns inherit a high mean, excess volatility and some time series predictability, again in line with the available evidence. Moreover, the risk-free rate is constant and low.

Another attractive feature of both narrow framing models is that they are able to generate excess volatility in the time series while still keeping the correlation of aggregate stock returns and consumption growth—a number that we also report in Table II—at realistically low levels. Some consumption-based models, such as Campbell and Cochrane (1999), are also able to generate excess volatility in aggregate returns. However, the extra volatility is driven by shocks to consumption growth, a mechanism that inevitably gives consumption growth and stock returns a counterfactually high correlation.

Of the two models, individual stock accounting may be the more successful one. Not only can it reproduce time series facts, but also a number of puzzling cross-sectional features of the data: the premium to value stocks and to stocks with poor prior returns, as well as the ability of certain multifactor models to capture the value premium. Portfolio accounting fails on the cross-sectional facts, and this is not simply a consequence of the particular parameter values we have used. The way the cross-sectional facts emerge under individual stock accounting is through a discount rate that is a function of a stock's own past performance. This is simply not a feature of portfolio accounting, whatever the parameters.

While individual stock accounting can be a useful device for understanding certain features of asset prices, we emphasize that it is at most a potential ingredient in an equilibrium model, and not a complete description of the facts. There are a number of dimensions on which it too, fails. For example, it predicts that the correlation between returns on different stocks is the same as the correlation between their cash flows. This can be seen in Table II, which reports the typical correlation between stocks to be $\omega_R = 0.23$, identical to the correlation between cash-flow shocks, $\omega_{ij} = 0.23$, listed in Table I. However, Vuolteenaho (1999) finds that shocks to expected returns on different stocks are actually much more correlated than cash-flow shocks, which immediately implies that stocks are more correlated with each other than are their underlying cash flows.²⁶

²⁶ Equivalently, the individual stock accounting model predicts that individual stock returns and aggregate stock returns will exhibit the same amount of excess volatility, while Vuolteenaho (1999) finds there to be much less excess volatility in individual stocks than in the aggregate stock market.

Moreover, while individual stock accounting does generate time series predictability, it does not generate enough: The R^2 in a time series regression of four-year cumulative aggregate stock returns on the lagged dividend yield is much smaller than the empirical value, reported by Fama and French (1988) to be over 20 percent.

One final reason why individual stock accounting may only be a partial explanation of the facts is that it also predicts that more volatile stocks will earn higher average returns, even though there is as yet little evidence of such an effect.²⁷ It is worth noting however, that individual stock accounting may be able to generate excess volatility in the time series as well as a value premium in the cross section *without* generating a large premium for idiosyncratic risk. The reason is that in our model, the first set of effects are generated by *changes* in the degree of loss aversion, while the price of volatility risk is determined by the average *level* of loss aversion. If investors' loss aversion changes over time without being too high on average, our model may be able to replicate the salient features of the data without producing a counterfactually high premium for volatility risk.

In summary, while individual stock accounting may offer a simple way of understanding a wide range of facts, it cannot be the complete story. A model that combines individual stock accounting with broader forms of accounting—portfolio accounting or even loss aversion over total wealth fluctuations—is likely to be better. Portfolio accounting introduces a common component in discount rate variation across stocks and therefore makes stocks move together more than their cash flows do. For reasons of tractability, we do not attempt an analysis of such a model in this paper.

A. Other Predictions

The fact that we study equilibrium returns under both individual stock accounting and portfolio accounting can also be useful for making additional predictions for future testing. If individual stock accounting is relatively more prevalent among individual investors as opposed to institutional investors, we would expect to see stocks held primarily by individuals—small stocks, for example—exhibit more of the features associated with individual stock accounting. In particular, by comparing the “Individual Accounting” and “Portfolio Accounting” columns in Table II, the specific prediction is that small stocks should have higher mean returns and more excess volatility than large stocks, should be more predictable in the time series and less correlated with each other, and that the value and De Bondt–Thaler premia should be stronger among small stocks.

Other predictions arise, if, over time, investors change the way they do their mental accounting. For example, the increased availability of mutual funds since the early 1980s may have caused a shift away from individual

²⁷ This is not to say that there is *no* evidence of such an effect: Using a different methodology from earlier studies, Lehmann (1990) does find some evidence of a residual risk premium.

Table V
The Risk of Value Strategies

The table reports the standard deviation, in simulated data, of a strategy which each year, sorts stocks by price–dividend ratio and buys the bottom decile (value stocks) and shorts the top decile (growth stocks). The calculation is done for the case when firm-level cash flows have a one-factor structure and for the case where they have a multifactor structure that allows for industry shocks. There are n stocks in total, and in the multifactor case, five industries.

Number of Stocks n	One-factor	Multifactor
100	18.2%	23.2%
500	11.4%	19.2%
1,000	9.7%	17.9%

stock accounting and towards portfolio accounting, since funds automatically prevent investors from worrying about individual stock fluctuations. Our analysis predicts that stocks that were once held directly but are now increasingly held indirectly through mutual funds should exhibit specific changes in pricing behavior. Among other predictions, such stocks should have higher price-to-fundamentals ratios and exhibit a lower cross-sectional value premium.

B. Exploiting Investors Who Frame Narrowly

In Section III.A, we introduced a one-factor cash-flow structure for stocks so as to simplify the calibration process as much as possible. If in reality, cash flows did indeed have a one-factor structure, it would be relatively straightforward for investors who do *not* frame narrowly to take advantage of investors who do, thus attenuating their effects. For example, they could buy a portfolio of stocks with price–dividend ratios in the lowest decile and short a portfolio of stocks with price–dividend ratios in the highest decile. This strategy earns the sizable value premium documented in Tables II and III, and if implemented with a very large number of stocks, becomes almost riskless: idiosyncratic risk gets washed away within the portfolio while the long-short position nets out the market factor. Table V provides some numbers: the column marked “One-factor” reports the standard deviation of the strategy just described for various values of n , the total number of stocks. As n increases, the standard deviation falls, dropping to as low as 9.7 percent for an n of 1,000. Coupling this standard deviation with the kind of value premia reported in Table II produces very attractive Sharpe ratios.

The problem with such a strategy in reality is that stock returns are driven not by one, but by many different factors, making it much harder to reduce the strategy’s risk, even with many stocks. To illustrate, the column titled “Multifactor” in Table V reports the standard deviation of the same long–short strategy, only this time computed for the more realistic cash-flow structure in equation (57). In particular, this cash-flow structure allows industry

factors to affect groups of stocks. Note that once we recognize these additional factors, the strategy becomes much more risky: going from 100 to 1,000 stocks only reduces risk by about a fifth, in contrast to the almost 50 percent reduction in risk in the “one-factor” case. Intuitively, many value stocks now belong to the same sector and hence co-move: There may have been bad news at the sector level, pushing all stocks in the sector down and leading investors to view them all as more risky, thus giving them low price–dividend ratios and making them all value stocks. Similarly, many growth stocks now belong to the same sector, and therefore also co-move. A long value, short growth strategy does not net out these industry factors, leaving it far from riskless. Daniel, Hirshleifer, and Subrahmanyam (2001) provide a similar discussion of this point in a related model.

There may be other reasons why exploiting investors who engage in narrow framing may be harder in reality than it appears in theory. It is likely to take many years of data before arbitrageurs can be statistically confident of the existence of a value premium, and even more time before they can convince themselves that it is not simply compensation for a lurking risk factor. Put simply, the effects of narrow framing may persist for a very long time before they can be detected and exploited.²⁸

C. Heterogeneity, Aggregation, and Trading

Even investors who do frame narrowly are likely to be heterogeneous in a number of ways: in the stocks that they hold or in the time at which they bought a particular stock. One limitation of our framework is the assumption that investors are completely homogeneous. This raises the question of whether the intuition of our models still holds once heterogeneity is recognized.

A full quantitative analysis of this issue is beyond the scope of this paper, but there is nonetheless reason to be hopeful that a more general model might deliver similar results. Consider the case of individual stock accounting. First, it is not at all clear that loss aversion will “wash out” in the aggregate. All investors holding a particular stock will find its fluctuations painful and will therefore require a high average return on the stock.

Moreover, if a stock goes down in value, all investors holding the stock will have their prior gains on the stock reduced, increasing their sensitivity to further losses regardless of when they bought into the stock or at what price. These investors now want to get rid of the stock: their selling pressure depresses the price, making the stock a value stock and giving it a higher expected return. Other investors who were not originally holding the stock can of course attenuate this effect: Since they did not experience the loss, they do not require such a high rate of return. However, these other investors will trade cautiously for the same reason that arbitrageurs in general

²⁸ After all, the value premium in our simulated data, while certainly strong, has been computed using 10,000 years of data. Arbitrageurs looking for exploitable anomalies have much less data at their disposal.

will trade cautiously: Common factors in value stocks makes buying these stocks risky. The price pressure will therefore be only partially absorbed, and the value premium will persist.

Similarly, if a stock rises in value, all investors holding the stock will accumulate larger prior gains on the stock, lowering their sensitivity to future losses. They will want to buy more of the stock, exerting a buying pressure that pushes up the price, making the stock a growth stock and giving it a lower expected return. This time, investors not already holding the stock can attenuate the effect by shorting, but a common factor in growth stocks together with restrictions on shorting by mutual funds suggest that the effect may survive.

Recent evidence on trading behavior may be consistent with the predictions of this more general way of thinking about pricing effects. Hvidkjaer (2000) finds strong selling pressure among stocks with poor prior performance and strong buying pressure among stocks with good prior performance. Moreover, if we think of “narrow framers” as being individual investors and the arbitrageurs as being mainly institutions, we would expect to see individuals selling to institutions in market troughs and institutions selling to individuals at market peaks. There is in fact evidence of such an effect: Cohen (1999) examines the long-term buying and selling patterns of individuals and institutions, and finds exactly the pattern we predict, namely individuals selling to institutions in market troughs and vice-versa at market peaks.

It is interesting to compare our prediction (and Cohen’s (1999) and Hvidkjaer’s (2000) results) with the findings of Odean (1998), who shows that individual investors are reluctant to sell stocks with short-term prior losses, preferring to sell prior short-term winner stocks.²⁹ The difference between this finding and our prediction may lie in the horizon over which investors come to terms with an investment loss. If a stock has experienced a short-term loss, the investor may not count this as a loss, preferring to hold on the stock in the hope of breaking even down the line. However, if the stock experiences a sustained, long-term drop, the investor may eventually decide that the investment has been a failure; he will accept the loss, view the stock as riskier, and then be ready to sell it.

V. Conclusion

A substantial body of experimental evidence suggests that loss aversion—the tendency to be more sensitive to losses than to gains—and narrow framing—the tendency to focus on narrowly defined gains and losses—play an important role in determining how people evaluate risky gambles. In this paper, we incorporate these ideas into an asset-pricing framework to see if they can shed light on the behavior of firm-level stock returns. In particular,

²⁹ This is often known as the disposition effect, and is also discussed by Shefrin and Statman (1985).

we consider two economies: one, labeled “individual stock accounting,” in which investors are loss averse over individual stock fluctuations, and another, labeled “portfolio accounting,” in which investors are loss averse only over portfolio fluctuations.

We find that both forms of narrow framing can explain certain aspects of the data. In both cases, the typical individual stock has a high mean, is excessively volatile, and is predictable in the time series. The aggregate stock return inherits these properties and is also only weakly correlated with consumption growth, while the risk-free rate is constant and low. Moreover, under individual stock accounting, there is a substantial value premium in the cross section, and this premium can to some extent be captured by the same kinds of multifactor models that have been successful in actual data.

Our paper suggests that using experimental evidence to refine the way we model investor preferences may be a promising avenue for further research. Nonetheless, there are alternative explanations under development for some of the phenomena discussed here, ranging from models firmly rooted in the consumption-based paradigm to models of learning and models of irrational beliefs. An important topic for further study is to clarify the distinct predictions of these various explanations so that additional testing might tell them apart.

Appendix

Proof of Proposition 1: Each agent’s optimization problem is

$$\max_{C_t, \{S_{i,t}\}} E \sum_{t=0}^{\infty} \left[\rho^t \frac{C_t^{1-\gamma}}{1-\gamma} + b_0 \bar{C}_t^{-\gamma} \rho^{t+1} \sum_{i=1}^n v(X_{i,t+1}, S_{i,t}, z_{i,t}) \right], \quad (\text{A1})$$

subject to the standard budget constraint

$$W_{t+1} = (W_t - C_t)R_f + \sum_{i=1}^n S_{i,t}(R_{i,t+1} - R_f), \quad (\text{A2})$$

where W_t denotes the agent’s preconsumption wealth at t .

The Euler equations for the optimization problem are

$$1 = \rho R_f E_t \left[\left(\frac{\bar{C}_{t+1}}{\bar{C}_t} \right)^{-\gamma} \right], \quad (\text{A3})$$

$$1 = \rho E_t \left[R_{i,t+1} \left(\frac{\bar{C}_{t+1}}{\bar{C}_t} \right)^{-\gamma} \right] + b_0 \rho E_t [\hat{v}(R_{i,t+1}, z_{i,t})], \quad \forall i. \quad (\text{A4})$$

These Euler equations are necessary conditions of optimality for the individual’s intertemporal problem. We now show that they are sufficient conditions, using a technique developed by Duffie and Skiadas (1994).

To simplify notation, let $u_t(C_t) = \rho^t C_t^{1-\gamma}/(1-\gamma)$ and $\bar{b}_t = \rho^{t+1} b_0 \bar{C}_t^{-\gamma}$. Assume that the strategy $(C^*, \{S_i^*\})$ satisfies the Euler equations

$$u'_t(C_t^*) = R_f E_t[u'_{t+1}(C_{t+1}^*)], \quad (\text{A5})$$

$$u'_t(C_t^*) = E_t[R_{i,t+1} u'_{t+1}(C_{t+1}^*)] + \bar{b}_t E_t[\hat{v}(R_{i,t+1}, z_{i,t})]. \quad (\text{A6})$$

Consider any alternative strategy $(C^* + \delta C, \{S_i^* + \delta S_i\})$ that satisfies the budget constraint

$$\delta W_{t+1} = (\delta W_t - \delta C_t) R_f + \sum_{i=1}^n \delta S_{i,t} (R_{i,t+1} - R_f). \quad (\text{A7})$$

The increase in expected utility from using the alternative strategy is

$$\begin{aligned} E \left[\sum_{t=0}^{\infty} \left[u_t(C_t^* + \delta C_t) - u_t(C_t^*) + \bar{b}_t \sum_{i=1}^n \delta S_{i,t} \hat{v}(R_{i,t+1}, z_{i,t}) \right] \right] \\ \leq \Delta \equiv E \left[\sum_{t=0}^{\infty} \left[u'_t(C_t^*) \delta C_t + \bar{b}_t \sum_{i=1}^n \delta S_{i,t} \hat{v}(R_{i,t+1}, z_{i,t}) \right] \right], \end{aligned} \quad (\text{A8})$$

where we have made use of the concavity of $u_t(\cdot)$ and the linearity of the prospect utility term with respect to $S_{i,t}$. It is therefore enough to show that $\Delta = 0$ under budget constraint (A7).

Multiplying equation (A7) by $u'_{t+1}(C_{t+1}^*)$ and applying Euler equations (A5) and (A6), we have

$$\begin{aligned} E_t \left[u'_t(C_t^*) \delta C_t + \bar{b}_t \sum_{i=1}^n \delta S_{i,t} \hat{v}(R_{i,t+1}, z_{i,t}) \right] \\ = u'_t(C_t^*) \delta W_t - E_t[u'_{t+1}(C_{t+1}^*) \delta W_{t+1}]. \end{aligned} \quad (\text{A9})$$

Summing up equation (A9) for all t and taking expectations, we have

$$\Delta = u_0(C_0^*) \delta W_0 - \lim_{T \rightarrow \infty} E[u'_T(C_T^*) \delta W_T]. \quad (\text{A10})$$

The budget constraint implies that $\delta W_0 = 0$. By requiring feasible strategies to use bounded units of financial securities, and with a unit of the risk-free security priced at one, we can show that the limiting term also goes to zero if our model parameters satisfy $\log \rho - \gamma g_c + g_i + 1/2(\gamma^2 \sigma_c^2 - 2\gamma \omega_{ci} \sigma_c \sigma_i + \sigma_i^2) < 0$, a condition that we already noted in footnote 12.³⁰ So $\Delta = 0$ for any feasible alternative to $(C^*, \{S_i^*\})$.

³⁰ The proof of this is available upon request.

We have therefore shown that any other budget feasible strategy cannot increase utility. The Euler equations are therefore necessary and sufficient conditions of optimality.

To close the proof, we show that if the risk-free gross interest rate is constant at R_f given by equation (21) and if the returns on any stock have a one-factor Markov structure given by equations (8), (19), and (20), with $f_i(\cdot)$ satisfying equation (22) for all $z_{i,t}$, then the agent's strategy of consuming his labor income and the dividend of each stock and holding the total supply of assets at each time t indeed satisfies the Euler equations.

The choice of risk-free rate in equation (21) allows us to satisfy Euler equation (A3), given equation (17). Given the assumed Markov structure of stock i 's returns, $z_{i,t+1}$ is determined by $z_{i,t}$ and $\varepsilon_{i,t+1}$. So we have

$$\begin{aligned} E_t \left[R_{i,t+1} \left(\frac{\bar{C}_{t+1}}{\bar{C}_t} \right)^{-\gamma} \right] &= E_t \left[\frac{1 + f_i(z_{i,t+1})}{f_i(z_{i,t})} e^{g_i + \sigma_i \varepsilon_{i,t+1}} e^{-\gamma(g_c + \sigma_c \eta_{t+1})} \right] \\ &= e^{g_i - \gamma g_c} E_t \left[E[e^{-\gamma \sigma_c \eta_{t+1}} | \varepsilon_{i,t+1}] e^{\sigma_i \varepsilon_{i,t+1}} \frac{1 + f_i(z_{i,t+1})}{f_i(z_{i,t})} \right] \\ &= e^{g_i - \gamma g_c + (1/2)\gamma^2 \sigma_c^2 (1 - \omega_{ci}^2)} E_t \left[\frac{1 + f_i(z_{i,t+1})}{f_i(z_{i,t})} e^{(\sigma_i - \gamma \omega_{ci} \sigma_c) \varepsilon_{i,t+1}} \right]. \end{aligned} \quad (\text{A11})$$

Applying this, we find that the strategy of consuming $\bar{C}_t$ and holding the supply of all securities satisfies Euler equation (A4) for all i . Q.E.D.

Proof of Proposition 2: Each agent's optimization problem is

$$\max_{C_t, \{S_{i,t}\}} E \sum_{t=0}^{\infty} \left[\rho^t \frac{C_t^{1-\gamma}}{1-\gamma} + b_0 \bar{C}_t^{-\gamma} \rho^{t+1} v(X_{t+1}, S_t, z_t) \right], \quad (\text{A12})$$

subject to the standard budget constraint

$$W_{t+1} = (W_t - C_t)R_f + \sum_{i=1}^n S_{i,t}(R_{i,t+1} - R_f), \quad (\text{A13})$$

where W_t denotes the agent's pre-consumption wealth at t .

The Euler equations for the optimization problem are

$$1 = \rho R_f E_t \left[\left(\frac{\bar{C}_{t+1}}{\bar{C}_t} \right)^{-\gamma} \right], \quad (\text{A14})$$

$$1 = \rho E_t \left[R_{i,t+1} \left(\frac{\bar{C}_{t+1}}{\bar{C}_t} \right)^{-\gamma} \right] + b_0 \rho E_t [\bar{v}(R_{i,t+1}, R_{t+1}, z_t)], \quad \forall i. \quad (\text{A15})$$

These Euler equations are necessary conditions of optimality for the individual's intertemporal problem. One can also show that they are sufficient conditions, using an identical argument to the one in the proof of Proposition 1. Summing up equation (A15) for all the stocks in the economy, we obtain the Euler equation for the market portfolio,

$$1 = \rho E_t \left[R_{t+1} \left(\frac{\bar{C}_{t+1}}{\bar{C}_t} \right)^{-\gamma} \right] + b_0 \rho E_t [\hat{v}(R_{t+1}, z_t)]. \quad (\text{A16})$$

To close the proof, we show that if the risk-free gross interest rate is constant at R_f given by equation (31), and if the returns on any stock have a one-factor Markov structure given by equations (15), (29), and (30), with $f_i(\cdot)$ satisfying equation (32) for all $z_{i,t}$, and if aggregate stock returns are given by equation (35) with $f(\cdot)$ satisfying equation (36) for all z_t , then the agent's strategy of consuming his labor income and the dividend of each stock and holding the total supply of assets at each time t indeed satisfies the Euler equations.

The choice of risk-free rate in equation (31) allows us to satisfy Euler equation (A14), given equation (17). A similar procedure to that used in the proof of Proposition 1 can be used to prove that the strategy of consuming $\bar{C}_t$ and holding the supply of all securities satisfies Euler equation (A16) under the conjectured Markov structure of market returns.

Turning now to Euler equation (A15), note that

$$E_t [e^{g_i + \sigma_i \varepsilon_{i,t+1}} | \varepsilon_{t+1}] = e^{g_i + (1/2)\sigma_i^2(1 - \omega_{ip}^2) + \sigma_i \omega_{ip} \varepsilon_{t+1}},$$

and

$$\begin{aligned} E [e^{-\gamma(g_c + \sigma_c \eta_{t+1})} e^{g_i + \sigma_i \varepsilon_{i,t+1}} | \varepsilon_{t+1}] \\ = e^{g_i - \gamma g_c + (1/2)\gamma^2 \sigma_c^2(1 - \omega_{cp}^2) + (1/2)\sigma_i^2(1 - \omega_{ip}^2) - \gamma \sigma_c \sigma_i (\omega_{ci} - \omega_{cp} \omega_{ip})} e^{(\sigma_i \omega_{ip} - \gamma \sigma_c \omega_{cp}) \varepsilon_{t+1}}. \end{aligned}$$

By iterated expectations, and using the above two equations and observing that z_{t+1} is determined by z_t and ε_{t+1} given the conjectured Markov structure of stock returns, we see that Euler equation (A15) reduces to equation (32) and is therefore satisfied by the strategy of consuming $\bar{C}_t$ and holding the supply of all stocks. Q.E.D.

REFERENCES

- Anderson, Robert, 1991, Nonstandard methods with applications in economics, in Werner Hildenbrand and Hugo Sonnenschein, eds.: *Handbook of Mathematical Economics* (North Holland, Amsterdam).
- Barberis, Nicholas, Ming Huang, and Tano Santos, 2001, Prospect theory and asset prices, *Quarterly Journal of Economics* 116, 1–53.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, 1998, A model of investor sentiment, *Journal of Financial Economics* 49, 307–343.

- Basu, Sanjoy, 1983, The relationship between earnings yield, market value and return for NYSE common stocks: Further evidence, *Journal of Financial Economics* 12, 129–156.
- Benartzi, Shlomo, and Richard Thaler, 1995, Myopic loss aversion and the equity premium puzzle, *Quarterly Journal of Economics* 110, 73–92.
- Benartzi, Shlomo, and Richard Thaler, 1999, Risk aversion or myopia? Choices in repeated gambles and retirement investments, *Management Science* 45, 364–381.
- Berk, Jonathan, Richard Green, and Vasant Naik, 1999, Optimal investment, growth options, and security returns, *Journal of Finance* 54, 1553–1607.
- Brennan, Michael, and Yihong Xia, 2001, Stock price volatility and the equity premium, *Journal of Monetary Economics*, forthcoming.
- Campbell, John Y., 2000, Asset prices, consumption and the business cycle, in John Taylor and Michael Woodford, eds.: *Handbook of Macroeconomics* (North Holland, Amsterdam).
- Campbell, John Y., and John Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.
- Campbell, John Y., and Robert Shiller, 1988, Stock prices, earnings, and expected dividends, *Journal of Finance* 43, 661–676.
- Cecchetti, Stephen, Pok-sang Lam, and Nelson Mark, 1990, Mean reversion in equilibrium asset prices, *American Economic Review* 80, 398–418.
- Cohen, Randolph, 1999, Asset allocation decisions of individuals and institutions, Working paper, Harvard Business School.
- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam, 2001, Covariance risk, mispricing, and the cross section of security returns, *Journal of Finance*, forthcoming.
- De Bondt, Werner, and Richard Thaler, 1985, Does the stock market overreact? *Journal of Finance* 40, 793–805.
- Duffie, Darrell, and Costis Skiadas, 1994, Continuous-time security pricing: A utility gradient approach, *Journal of Mathematical Economics* 23, 107–131.
- Epstein, Larry, and Stanley Zin, 1990, First-order risk aversion and the equity premium puzzle, *Journal of Monetary Economics* 26, 387–407.
- Fama, Eugene, and Kenneth French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3–27.
- Fama, Eugene, and Kenneth French, 1992, The cross section of expected stock returns, *Journal of Finance* 47, 427–465.
- Fama, Eugene, and Kenneth French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Gneezy, Uri, and Jan Potters, 1997, An experiment on risk taking and evaluation periods, *Quarterly Journal of Economics* 112, 631–645.
- Gomes, Joao, Leonid Kogan, and Lu Zhang, 2001, Equilibrium cross-section of returns, Working paper, University of Pennsylvania.
- Green, Edward, 1989, Individual-level randomness in a nonatomic population, Working paper, University of Pittsburgh.
- Gul, Faruk, 1991, A theory of disappointment aversion, *Econometrica* 59, 667–686.
- Hong, Harrison, and Jeremy Stein, 1999, A unified theory of underreaction, momentum trading, and overreaction in asset markets, *Journal of Finance* 54, 2143–2184.
- Hvidkjaer, Soeren, 2000, A trade-based analysis of momentum, Working paper, Cornell University.
- Kahneman, Daniel, and Dan Lovallo, 1993, Timid choices and bold forecasts: A cognitive perspective on risk taking, *Management Science* 39, 17–31.
- Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263–291.
- Lehmann, Bruce, 1990, Residual risk revisited, *Journal of Econometrics* 45, 71–97.
- LeRoy, Stephen, and Richard Porter, 1981, The present value relation: Tests based on variance bounds, *Econometrica* 49, 555–577.
- Lucas, Robert, 1978, Asset prices in an exchange economy, *Econometrica* 46, 1419–1446.
- Odean, Terrance, 1998, Are investors reluctant to realize their losses? *Journal of Finance* 53, 1775–1798.

- Rabin, Matthew, 2000, Risk aversion and expected utility theory: A calibration theorem, *Econometrica* 68, 1281–1292.
- Rabin, Matthew, and Richard Thaler, 2000, Anomalies: Risk aversion, Working paper, University of Chicago.
- Read, Daniel, George Loewenstein, and Matthew Rabin, 1999, Choice bracketing, *Journal of Risk and Uncertainty* 19, 171–197.
- Redelmeier, Daniel, and Amos Tversky, 1992, On the framing of multiple prospects, *Psychological Science* 3, 191–193.
- Rosenberg, Barr, Kenneth Reid, and Ronald Lanstein, 1985, Persuasive evidence of market inefficiency, *Journal of Portfolio Management* 11, 9–17.
- Segal, Uzi, and Avia Spivak, 1990, First-order versus second-order risk aversion, *Journal of Economic Theory* 51, 111–125.
- Shefrin, Hershey, and Meir Statman, 1985, The disposition to sell winners too early and ride losers too long: Theory and evidence, *Journal of Finance* 40, 777–790.
- Shiller, Robert, 1981, Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71, 421–436.
- Thaler, Richard, 1980, Toward a positive theory of consumer choice, *Journal of Economic Behavior and Organization* 1, 39–60.
- Thaler, Richard, and Eric Johnson, 1990, Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice, *Management Science* 36, 643–660.
- Thaler, Richard, Amos Tversky, Daniel Kahneman, and Alan Schwarz, 1997, The effect of myopia and loss aversion on risk taking: An experimental test, *Quarterly Journal of Economics* 112, 647–661.
- Tversky, Amos, and Daniel Kahneman, 1981, The framing of decisions and the psychology of choice, *Science* 211, 453–458.
- Tversky, Amos, and Daniel Kahneman, 1992, Advances in prospect theory: Cumulative representation of uncertainty, *Journal of Risk and Uncertainty* 5, 297–323.
- Vuolteenaho, Tuomo, 1999, What drives firm level stock returns?, Working paper, University of Chicago.

Discussion

M. J. BRENNAN*

THIS IS A STIMULATING and provocative paper. It not only demands our professional respect for its modeling skill, but also forces us to reflect on the nature of financial modeling, and on how far from rationality we should be willing to depart in our assumptions about investor behavior. Conventions here have clearly changed over time.

In the years that followed the development of the Capital Asset Pricing Model in the mid-1960s, before the advent of formal behavioral models of financial markets, convention required modelers to remain for the most part

* Goldyne and Irwin Hearsh Professor, University of California, Los Angeles, and Professor of Finance, London Business School.