THE JOURNAL OF FINANCE • VOL. LXIV, NO. 2 • APRIL 2009

What Drives the Disposition Effect? An Analysis of a Long-Standing Preference-Based Explanation

NICHOLAS BARBERIS and WEI XIONG*

ABSTRACT

We investigate whether prospect theory preferences can predict a disposition effect. We consider two implementations of prospect theory: in one case, preferences are defined over *annual* gains and losses; in the other, they are defined over *realized* gains and losses. Surprisingly, the annual gain/loss model often fails to predict a disposition effect. The realized gain/loss model, however, predicts a disposition effect more reliably. Utility from realized gains and losses may therefore be a useful way of thinking about certain aspects of individual investor trading.

One of the most robust facts about the trading of individual investors is the "disposition effect": when an individual investor sells a stock in his portfolio, he has a greater propensity to sell a stock that has gone *up* in value since purchase than one that has gone down. The effect has been documented in all the available large databases of individual investor trading activity and has been linked to important pricing phenomena such as post-earnings announcement drift and stock-level momentum. Disposition effects have also been uncovered in other settings—in the real estate market, for example, and in the exercise of executive stock options.[1]

While the disposition effect is a fundamental feature of trading, its underlying cause remains unclear. *Why* do individual investors have a greater propensity to

*Barberis is at the Yale School of Management. Xiong is at the Department of Economics, Princeton University. We are grateful to Campbell Harvey, an associate editor, two anonymous referees, John Campbell, Francisco Gomes, Mark Grinblatt, Bing Han, Xuedong He, Daniel Kahneman, George Loewenstein, Cade Massey, Mark Salmon, Jeremy Stein, Richard Thaler, Mark Westerfield, and seminar participants at Carnegie-Mellon University, Duke University, Imperial College, Northwestern University, Princeton University, Rutgers University, the Stockholm Institute for Financial Research, the University of California at Berkeley, the University of North Carolina, the University of Texas, the University of Warwick, the University of Wisconsin, Yale University, the Utah Winter Finance Conference, the Western Finance Association, and the NBER for helpful feedback.

[1] Odean (1998), Grinblatt and Keloharju (2001), and Feng and Seasholes (2005) document the disposition effect for individual investors in the U.S., Finland, and China, respectively. Frazzini (2006) finds a disposition effect in the trading of U.S. mutual fund managers, albeit weaker than that for individual investors. Grinblatt and Han (2005) and Frazzini (2006) produce evidence linking the disposition effect to momentum and post-earnings announcement drift. Genesove and Mayer (2001) and Heath, Huddart, and Lang (1999) uncover disposition effects in the real estate market and in the exercise of executive stock options.

sell stocks trading at a paper gain rather than those trading at a paper loss? In a careful study of the disposition effect, Odean (1998) shows that the most obvious potential explanations—explanations based on informed trading, rebalancing, or transaction costs—fail to capture important features of the data.

Given the difficulties faced by standard hypotheses, an alternative view based on prospect theory has gained favor. Prospect theory, a prominent theory of decision-making under risk proposed by Kahneman and Tversky (1979) and extended by Tversky and Kahneman (1992), posits that people evaluate gambles by thinking about gains and losses, not final wealth levels, and that they process these gains and losses using a value function that is concave for gains and convex for losses. The value function is designed to capture the experimental finding that people tend to be risk-averse over moderate-probability gains (they typically prefer a certain \$100 to a 50:50 bet to win \$0 or \$200), but tend to be risk-*seeking* over moderate-probability losses (they typically prefer a 50:50 bet to lose \$0 or \$200 to a certain loss of \$100).

Prospect theory is potentially a useful ingredient in a model of the disposition effect. If an investor is holding a stock that has risen in value since purchase, he may think of the stock as trading at a gain. If he is risk-averse over gains, he may then be inclined to sell the stock. Similarly, if he is risk-seeking over losses, he may be inclined to hold on to a stock that has gone down in value.

While prospect theory does seem to offer a promising framework for thinking about the disposition effect, the link has almost always been discussed in informal terms. This leaves a number of questions unanswered: Can a link between prospect theory and the disposition effect be formalized in a rigorous model? Under what conditions does prospect theory predict a disposition effect? What other predictions does prospect theory make about trading activity? To answer these questions, some formal modeling is needed.

In this paper, we take up this task, and study the trading behavior of an investor with prospect theory preferences. We consider two implementations of prospect theory. The first implementation, which is the focus of Section II, applies prospect theory to annual stock-level trading profits. Specifically, we consider an investor who, at the beginning of the year, buys shares of a stock. Over the course of the year, he trades the stock, and, at the end of the year, receives prospect theory utility based on his trading profit. The year is divided into $T \geq 2$ trading periods. We use the prospect theory value function proposed by Tversky and Kahneman (1992). For much of the analysis, we also use the preference parameters these authors estimate from experimental data.

For any T, we obtain an analytical solution for the optimal trading strategy. This allows us to simulate artificial data on how prospect theory investors would trade over time, and then to check, using Odean's (1998) methodology, whether prospect theory predicts a disposition effect. We pay particular attention to how the results depend on the expected stock return μ and the number of trading periods T.

Our analysis leads to a surprising finding. While for some values of μ and T this implementation of prospect theory does predict a disposition effect, for many other values of μ and T it predicts the *opposite* of the disposition effect, namely that investors will be more inclined to sell stocks with prior *losses* than stocks with prior gains. We explain the intuition for this result in Section II.B by way of a detailed example.

In Section III, we consider a second implementation of prospect theory, one in which prospect theory is defined over *realized* gains and losses. In this case, if an investor buys some shares of a stock at the start of the year, and then, a few months later, sells some of the shares, he receives a jolt of prospect theory utility right then, at the moment of sale, where the utility term depends on the size of the realized gain or loss. We find that this implementation leads more readily to a disposition effect, although even here, we occasionally observe the opposite of the disposition effect.

Of the two implementations, the second, which applies prospect theory to realized gains and losses, represents a more significant departure from the standard finance paradigm: It assumes not only prospect theory, but also that investor preferences distinguish between paper and realized gains. Our analysis shows that, even if this implementation is more radical, it deserves to be taken seriously: It predicts a disposition effect reliably, while a more standard model—one that applies prospect theory to *annual* gains and losses—does not.

In Section I, we review the evidence on the disposition effect and the elements of prospect theory. In Section II, we analyze trading behavior in a model that applies prospect theory to annual stock-level trading profits. In Section III, we consider an alternative implementation in which prospect theory is defined over realized gains and losses. Section IV discusses related research and other applications, and Section V concludes.

I. The Disposition Effect: Evidence and Interpretation

Odean (1998) analyzes the trading activity over the 1987 to 1993 period of 10,000 households with accounts at a large discount brokerage firm. He finds that, when an investor in his sample sells shares, he has a greater propensity to sell shares of a stock that has *risen* in value since purchase rather than of one that has fallen in value. Specifically, for any day on which an investor in the sample sells shares of a stock, each stock in his portfolio on that day is placed into one of four categories. For every stock in the investor's portfolio on that day that is *sold,* a "realized gain" is counted if the stock price exceeds the average price at which the shares were purchased, and a "realized loss" is counted otherwise. For every stock in the investor's portfolio on that day that is *not* sold, a "paper gain" is counted if the stock price exceeds the average price at which the shares were purchased, and a "paper loss" is counted otherwise. From the total number of realized gains and paper gains across all accounts over the entire sample, Odean (1998) computes the proportion of gains realized (PGR):

$$\text{PGR} = \frac{\text{no. of realized gains}}{\text{no. of realized gains} + \text{no. of paper gains}}. \tag{1}$$

In words, PGR computes the number of gains that were realized as a fraction of the total number of gains that could have been realized. A similar ratio,

$$\text{PLR} = \frac{\text{no. of realized losses}}{\text{no. of realized losses} + \text{no. of paper losses}}, \tag{2}$$

is computed for losses. The disposition effect is the empirical fact that PGR is significantly greater than PLR. Odean (1998) reports PGR = 0.148 and PLR = 0.098.

Robust as this effect is, its cause remains unclear: The most obvious potential explanations fail to capture important features of the data. Perhaps the most obvious hypothesis of all is the information hypothesis, namely that investors sell stocks with paper gains because they have private information that these stocks will subsequently do poorly, and hold on to stocks with paper losses because they have private information that these stocks will rebound. This hypothesis is refuted, however, by Odean's (1998) finding that the average return of prior winners that investors sell is 3.4% *higher*, over the next year, than the average return of the prior losers they hold on to.

Tax considerations also fail to shed light on the disposition effect: Such considerations predict a greater propensity to sell stocks with paper *losses* because the losses thus realized can be used to offset taxable gains in other assets.[2]

Odean (1998) also casts doubt on the hypothesis that the disposition effect is nothing more than portfolio rebalancing of the kind predicted by a model with power utility preferences and i.i.d. returns. He does so by showing that the disposition effect remains strong even when the sample is restricted to sales of investors' *entire* holdings of a stock. If rebalancing occurs at all, it is more likely to manifest itself as a *partial* reduction of a stock position that has risen in value, rather than as a sale of the entire position. Another difficulty with the rebalancing view is that, since, under this view, rebalancing is the "smart" thing to do, the disposition effect should be stronger for more sophisticated investors. In actuality, however, it is the *less* sophisticated investors who exhibit the disposition effect more strongly (Dhar and Zhu (2006)).

Given the difficulties faced by these standard hypotheses, an alternative view of the disposition effect based on Kahneman and Tversky's (1979) prospect theory has gained favor. We now briefly review the main features of prospect theory.

A. Prospect Theory

Consider the gamble

$$(x, p; y, q),$$

to be read as "gain x with probability p and y with probability q, independent of other risks," where $x \leq 0 \leq y$ or $y \leq 0 \leq x$, and where $p + q = 1$. In the expected

[2] Odean (1998) finds that, in one month of the year, December, PLR exceeds PGR. This suggests that tax factors play a larger role as the end of the tax year approaches.

utility framework, an agent with utility function $U(\cdot)$ evaluates this risk by computing

$$pU(W+x)+qU(W+y), \tag{3}$$

where W is his current wealth. By contrast, in the framework of prospect theory, the agent assigns the gamble the value

$$\pi(p)v(x)+\pi(q)v(y), \tag{4}$$

where $v(\cdot)$ and $\pi(\cdot)$ are known as the value function and the probability weighting function, respectively. These functions satisfy $v(0)=0$, $\pi(0)=0$, and $\pi(1)=1$.

There are four important differences between expressions (3) and (4). First, the carriers of value in prospect theory are gains and losses, not final wealth levels: The argument of $v(\cdot)$ in (4) is x, not $W+x$. Second, while $U(\cdot)$ is typically concave everywhere, $v(\cdot)$ is concave only over gains; over losses, it is convex: People tend to be risk-averse over moderate-probability gains but risk-seeking over moderate-probability losses. Third, while $U(\cdot)$ is typically differentiable everywhere, $v(\cdot)$ has a kink at the origin, so that the agent is more sensitive to losses, even small losses, than to gains of the same magnitude. This feature, known as loss aversion, is inferred from the widespread aversion to gambles such as a 50:50 bet to win \$110 or lose \$100. Finally, the prospect theory agent does not use objective probabilities, but rather, transformed probabilities obtained from objective probabilities via the weighting function $\pi(\cdot)$. The primary effect of this weighting function is to overweight low probabilities, a feature that parsimoniously captures the simultaneous demand many individuals have for both lottery tickets and insurance.

Tversky and Kahneman (1992) propose a specific form for the value function, namely

$$v(x)=\begin{cases} x^{\alpha} & \text{for} \quad x\geq 0 \\ -\lambda(-x)^{\beta} & \phantom{\text{for}} \quad x<0 \end{cases}. \tag{5}$$

For $\alpha, \beta \in (0,1)$ and $\lambda > 1$, this function is indeed concave over gains and convex over losses, and does indeed exhibit a greater sensitivity to losses than to gains. Using experimental data, Tversky and Kahneman (1992) estimate $\alpha=\beta=0.88$ and $\lambda=2.25$. Since the estimated values of α and β are the same, we will work, from this point on, with the specification

$$v(x)=\begin{cases} x^{\alpha} & \text{for} \quad x\geq 0 \\ -\lambda(-x)^{\alpha} & \phantom{\text{for}} \quad x<0 \end{cases}, \quad 0<\alpha<1,\ \lambda>1. \tag{6}$$

Figure 1 plots the function in equation (6) for $\alpha=0.88$ and $\lambda=2.25$. An α of 0.88 means that the function is only mildly concave for gains and only mildly

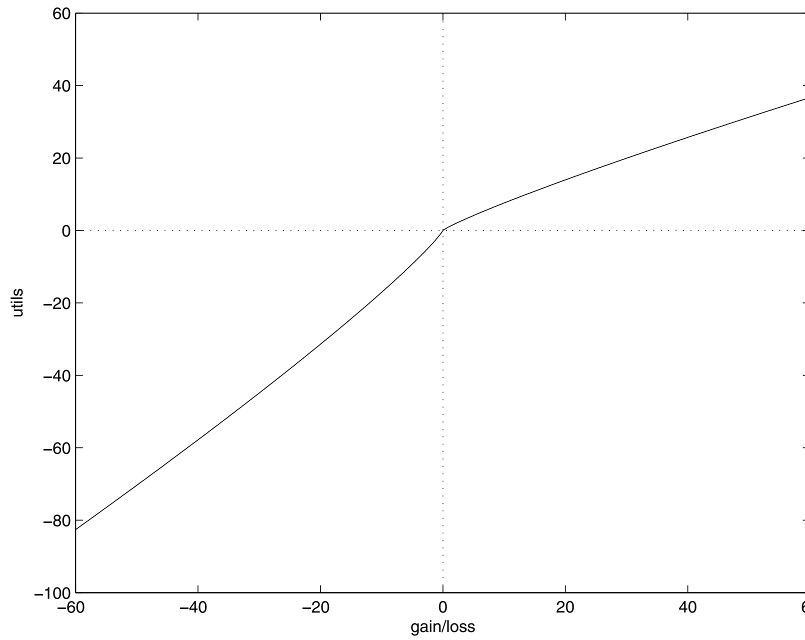

Figure 1. The prospect theory value function. The figure plots the prospect theory value function $v(\cdot)$ proposed by Tversky and Kahneman (1992). In the region of gains, $v(x) = x^{\alpha}$, and in the region of losses, $v(x) = -\lambda(-x)^{\alpha}$, where $\alpha = 0.88$ and $\lambda = 2.25$. The graph shows that, for these parameter values, the concavity in the region of gains and the convexity in the region of losses are both mild, while the kink at the origin is sharp.

convex for losses, while a λ of 2.25 implies much greater sensitivity to losses than to gains. This will be important in what follows.[3]

As noted in the introduction, prospect theory is potentially a useful ingredient in a model of the disposition effect. In the past, however, the link has almost always been discussed in informal terms. In Sections II and III, we investigate more formally whether, and under what conditions, prospect theory predicts a disposition effect.

II. A Model That Applies Prospect Theory to Annual Trading Profits

We consider a portfolio choice setting with $T+1$ dates, $t = 0, 1, \ldots, T$. There are two assets: a risk-free asset, which earns a gross return of $R_f \geq 1$ in each period, and a risky asset, which we think of as an individual stock. The price

[3] Strictly speaking, the function in equation (6) does not have a kink at the origin: $v'(x) \to \infty$ as $x \to 0$ from above or below. However, for $\lambda > 1$, it does satisfy $v(x) < -v(-x)$ for $x > 0$. In this sense, the agent is more sensitive to losses than to gains.

of the stock at time t is P_t. Its gross return from t to $t+1$, $R_{t,t+1}$, is distributed according to

$$R_{t,t+1} = \begin{cases} R_u > R_f & \text{with probability } \pi \\ R_d < R_f & \text{with probability } 1-\pi \end{cases}, \quad \text{i.i.d. across periods,} \tag{7}$$

so that the stock price evolves along a binomial tree. We assume

$$\pi R_u + (1-\pi)R_d > R_f, \tag{8}$$

so that the expected stock return exceeds the risk-free rate. When we calibrate the model in Section II.A, we take the interval from 0 to T to be a year.

We study the trading behavior of an investor with prospect theory preferences, who, in particular, uses the value function $v(\cdot)$ in equation (6). The argument of $v(\cdot)$ is the investor's "gain" or "loss". Defining the gain or loss is an important step.

In the stock market context, a gain or loss can be defined in a number of ways. We consider two possibilities. In Section II, we take the gain or loss to be the profit from trading the stock over the year-long interval between time 0 and time T. We refer to this as the "annual gain/loss" implementation of prospect theory. A similar implementation, one in which prospect theory is defined over annual gains and losses, has been used with some success in earlier applications of prospect theory to phenomena such as the equity premium and the low average return on IPOs (Benartzi and Thaler (1995), Barberis and Huang (2008)).

In Section III, we consider a second implementation in which prospect theory is defined over *realized* gains and losses: If the investor sells some shares at time t, $0 < t \leq T$, he receives a jolt of prospect theory utility right then, at time t, where the argument of the utility function is the size of the realized gain or loss.

Of the two implementations, the annual gain/loss model is closer to the standard finance paradigm: The realized gain/loss model appeals not only to prospect theory but also to a distinction between realized and paper gains, a distinction that finance models do not normally make when specifying preferences. There is a tradition in economic modeling that departures from the standard model are made incrementally, so that we can understand which assumptions are truly necessary in order to explain the facts. We follow this tradition by studying the annual gain/loss model first, and only then turning to the realized gain/loss model.

In Section II, then, we take the gain or loss to be the profit from trading the stock over the interval from 0 to T. One possible candidate for the argument of the value function $v(\cdot)$ is therefore $W_T - W_0$, where W_t is the investor's wealth at time t: After all, $W_T - W_0$ is, quite literally, the profit from trading between 0 and T. In our analysis, we actually define the gain or loss to be

$$\triangle W_T \equiv W_T - W_0 R_f^T. \tag{9}$$

This is the investor's trading profit over the interval from 0 to T relative to the profit he could have earned by investing in the risk-free asset. This definition is more tractable and may also be more plausible: The investor may only consider his trading a success if it earns him more than just the compounded risk-free return. We refer to $W_0 R_f^T$ as the "reference" level of wealth, so that the gain or loss is final wealth minus this reference wealth level.

For simplicity, we ignore probability weighting, so that the investor uses objective rather than transformed probabilities. The primary effect of probability weighting is to overweight low probabilities; it therefore has its biggest impact on securities with highly skewed returns (Barberis and Huang (2008)). Since most stocks are not highly skewed, we do not expect probability weighting to be central to the link between prospect theory and the disposition effect. We discuss this issue further in Section IV.[4]

At each date from $t = 0$ to $t = T - 1$, the investor must decide how to split his wealth between the risk-free asset and the risky asset. If x_t is the number of shares of the risky asset he holds at time t, his decision problem is

$$\max_{x_0, x_1, \ldots, x_{T-1}} E[v(\Delta W_T)] = E\big[v\big(W_T - W_0 R_f^T\big)\big], \tag{10}$$

where $v(\cdot)$ is defined in equation (6), subject to the budget constraint

$$\begin{aligned} W_t &= (W_{t-1} - x_{t-1}P_{t-1})R_f + x_{t-1}P_{t-1}R_{t-1,t} \\ &= W_{t-1}R_f + x_{t-1}P_{t-1}(R_{t-1,t} - R_f), \quad t = 1, \ldots, T, \end{aligned} \tag{11}$$

and a nonnegativity of wealth constraint

$$W_T \geq 0. \tag{12}$$

The problem in (10) to (12) assumes just one risky asset. However, under two conditions, its solution also describes optimal trading in a multi-stock setting. The first condition is that the investor engages in what is sometimes called "narrow framing" or "mental accounting," so that, even if he trades several stocks, he gets utility separately from the annual trading profit on each one. This assumption is always present in the informal arguments that have been used to link prospect theory and the disposition effect and we adopt it here, too. The second condition is that we reinterpret W_0 as the maximum amount the investor is willing to lose from trading any one of his stocks. Under these conditions, the investor's trading strategy for each stock is independent of his other holdings and is therefore given by the solution to (10) to (12).

A. The Optimal Trading Strategy

The problem in (10) to (12) can be solved analytically for any number of trading periods T. To obtain the solution, we use the insight of Cox and Huang (1989) who demonstrate that, when markets are complete, an investor's

[4] At the risk of causing confusion, we have used the notation $\pi(\cdot)$ for the probability weighting function that forms part of prospect theory and π for the probability of a good stock return. The function $\pi(\cdot)$ will not appear again in the paper; the variable π will.

dynamic optimization problem can be rewritten as a *static* problem in which the investor directly chooses his wealth in the different possible states at the final date. An optimal trading strategy is one that generates these optimal wealth allocations. In a complete market, such a trading strategy always exists.

To implement this technique in our context, some notation will be helpful. In our model, the price of the risky asset evolves along a binomial tree. At date t, there are $t+1$ nodes in the tree, $j = 1, 2, \ldots, t+1$, where $j = 1$ corresponds to the highest node in the tree at that date and $j = t+1$ to the lowest. The price of the risky asset in node j at time t, $P_{t,j}$, is $P_0 R_u^{t-j+1} R_d^{j-1}$.

We denote the optimal share allocation in node j at time t by $x_{t,j}$, the optimal wealth in that node by $W_{t,j}$, and the ex-ante probability of reaching that node by $\pi_{t,j}$, so that

$$\sum_{j=1}^{t+1} \pi_{t,j} = 1. \tag{13}$$

If $p_{t,j}$ is the time 0 price of a contingent claim that pays \$1 if the stock price reaches node j at time t, the state price density for that node is

$$q_{t,j} = \frac{p_{t,j}}{\pi_{t,j}}. \tag{14}$$

The state price density is linked to the risk-free rate by

$$\sum_{j=1}^{t+1} \pi_{t,j} q_{t,j} = \frac{1}{(R_f)^t}. \tag{15}$$

With this notation in hand, we apply Cox and Huang's (1989) insight and rewrite the problem in (10) to (12) as

$$\max_{\{W_{T,j}\}_{j=1,\ldots,T+1}} \sum_{j=1}^{T+1} \pi_{T,j} v\left(W_{T,j} - W_0 R_f^T\right), \tag{16}$$

subject to the budget constraint

$$\sum_{j=1}^{T+1} \pi_{T,j} q_{T,j} W_{T,j} = W_0 \tag{17}$$

and a nonnegativity of wealth constraint

$$W_{T,j} \geq 0, \quad j = 1, \ldots, T+1. \tag{18}$$

This static problem can be solved using the Lagrange multiplier method. We present the solution in Proposition 1 below. For simplicity, the proposition assumes $\pi = \frac{1}{2}$, so that, in each period, a good stock return and a poor stock return are equally likely. Under this assumption,

$$\pi_{t,j} = \frac{t!2^{-t}}{(t-j+1)!(j-1)!}. \tag{19}$$

In the proof of Proposition 1, we also show that, under this assumption,

$$q_{t,j} = q_u^{t-j+1} q_d^{j-1}, \tag{20}$$

where

$$q_u = \frac{2(R_f - R_d)}{R_f(R_u - R_d)}, \quad q_d = \frac{2(R_u - R_f)}{R_f(R_u - R_d)}, \tag{21}$$

so that the state price density increases as we go down the $t+1$ nodes at date t.

Proposition 1: *For $\pi = \frac{1}{2}$, the optimal wealth allocations $W_{t,j}$ and optimal share holdings of the risky asset $x_{t,j}$ can be obtained as follows. Let*

$$V^* = \max_{k\in\{1,\ldots,T\}} \left[\left(\sum_{l=1}^{k} q_{T,l}^{-\frac{\alpha}{1-\alpha}} \pi_{T,l} \right)^{1-\alpha} \left(\sum_{l=k+1}^{T+1} q_{T,l}\pi_{T,l} \right)^{\alpha} - \lambda \sum_{l=k+1}^{T+1} \pi_{T,l} \right], \tag{22}$$

and let k^* be the $k \in \{1, \ldots, T\}$ at which the maximum in (22) is attained.

Then, the optimal wealth allocation $W_{T,j}$ in node j at final date T is given by

$$W_{T,j} = \begin{cases} W_0 R_f^T \left[1 + q_{T,j}^{-\frac{1}{1-\alpha}} \dfrac{\sum_{l=k^*+1}^{T+1} q_{T,l}\pi_{T,l}}{\sum_{l=1}^{k^*} q_{T,l}^{-\frac{\alpha}{1-\alpha}} \pi_{T,l}} \right] & \text{if } j \le k^* \\ 0 & \text{if } j > k^* \end{cases} \tag{23}$$

if $V^* > 0$, and by

$$W_{T,j} = W_0 R_f^T, \quad j = 1, \ldots, T+1, \tag{24}$$

if $V^* \le 0$. The optimal share holdings $x_{t,j}$ are given by

$$x_{t,j} = \frac{W_{t+1,j} - W_{t+1,j+1}}{P_0\left(R_u^{t-j+2} R_d^{j-1} - R_u^{t-j+1} R_d^{j}\right)}, \quad 0 \le t \le T-1, \ 1 \le j \le t+1, \tag{25}$$

where the intermediate wealth allocations can be computed by working backwards from date T using

$$W_{t,j} = \frac{\frac{1}{2}W_{t+1,j}q_{t+1,j} + \frac{1}{2}W_{t+1,j+1}q_{t+1,j+1}}{q_{t,j}},$$
$$0 \le t \le T-1,\ 1 \le j \le t+1. \qquad (26)$$

Proof: See the Appendix.

Before analyzing the optimal share holdings $x_{t,j}$, we note some features of the optimal date T wealth allocations $W_{T,j}$ in (23) and (24). We find that the investor's optimal policy is either to choose an allocation equal to the reference wealth level $W_0R_f^T$ in all date T nodes, as in (24); or, as in (23), to use a "threshold" strategy in which, for some $k^* : 1 \le k^* \le T$, he allocates a wealth level greater than the reference level $W_0R_f^T$ to the k^* date T nodes with the lowest state price densities—in other words, the k^* date T nodes with the highest risky asset prices—and a wealth level of zero to the remaining date T nodes. To find the best threshold strategy, equation (22) maximizes the investor's utility across the T possible values of k^*. If the best threshold strategy offers nonpositive utility, that is, if $V^* \le 0$, which occurs when the expected risky asset return is low, then the investor does not use a threshold strategy and instead chooses a wealth level of $W_0R_f^T$ in all final date nodes; otherwise, he adopts the best threshold strategy.

The results in Proposition 1 are similar to those of Berkelaar, Kouwenberg, and Post (2004) who solve the continuous-time analog of (10) to (12), also using the Cox and Huang (1989) technique. In their model, the investor trades continuously from time 0 to time T' and, at time T', derives prospect theory utility from the difference between time T' wealth and a reference wealth level. As in our model, probability weighting is ignored. The continuous-time and discrete-time solutions are similar: In both cases, so long as the investor is willing to buy the risky asset at time 0, his optimal wealth at time T' is either zero or some amount that exceeds the reference wealth level.

In this paper, we use a discrete-time framework because we want to be able to vary how often the investor can change his share holdings and hence to see whether the link between prospect theory and the disposition effect depends on trading frequency. Berkelaar, Kouwenberg, and Post (2004) do not discuss the disposition effect; their focus is instead on how the investor's time 0 allocation depends on the variable T'.

We now illustrate Proposition 1 with an example. We set the initial price of the risky asset to $P_0 = 40$, the investor's initial wealth to $W_0 = 40$, the gross risk-free rate to $R_f = 1$, the number of periods to $T = 4$, and the preference parameters to $(\alpha, \lambda) = (0.88, 2.25)$, the values estimated by Tversky and Kahneman (1992) from experimental data.

We also need to assign values to R_u and R_d. To do this, we take the interval from $t = 0$ to $t = T$ to be a fixed length of time, namely a year. We choose plausible values for the *annual* gross expected return μ and standard deviation σ of the risky asset and then, for any T, back out the implied values of R_u and R_d. For $\pi = \frac{1}{2}$, R_u and R_d are related to μ and σ by

$$\left(\frac{R_u + R_d}{2}\right)^T = \mu, \quad \left(\frac{R_u^2 + R_d^2}{2}\right)^T = \mu^2 + \sigma^2, \tag{27}$$

which imply

$$R_u = \mu^{\frac{1}{T}} + \sqrt{(\mu^2 + \sigma^2)^{\frac{1}{T}} - (\mu^2)^{\frac{1}{T}}} \tag{28}$$

$$R_d = \mu^{\frac{1}{T}} - \sqrt{(\mu^2 + \sigma^2)^{\frac{1}{T}} - (\mu^2)^{\frac{1}{T}}}. \tag{29}$$

In our example, we set $(\mu, \sigma) = (1.1, 0.3)$, which, from (28) and (29), corresponds to $(R_u, R_d) = (1.16, 0.89)$.

For these parameter values, the top-left panel in Table I shows the binomial tree for the price of the risky asset. The top-right panel reports the state price density at each node in the tree, computed using equations (20) and (21). The bottom-left and bottom-right panels report optimal share holdings and optimal wealth allocations at each node, respectively.

The right-most column in the bottom-right panel illustrates one of the results in Proposition 1: The wealth allocation at the final date is either zero or a positive amount that exceeds the reference wealth level of \$40. Meanwhile, the optimal share holdings in the bottom-left panel provide an early hint of the results to come. If anything, the investor has a greater propensity to sell shares after a *drop* in the stock price rather than after a rise. This is the opposite of the disposition effect.[5]

We now investigate more carefully whether the annual gain/loss implementation of prospect theory predicts a disposition effect. In brief, we use Proposition 1 to simulate an artificial data set of how prospect theory investors would trade over time. We then apply Odean's (1998) methodology to see if, in the simulated data, investors exhibit a disposition effect.

Odean's (1998) data cover 10,000 households. We therefore generate trading data for 10,000 prospect theory investors, each of whom holds N_S stocks. For each investor, we use the binomial distribution in (7) with $\pi = \frac{1}{2}$ to simulate a

[5] Beyond the lack of an obvious disposition effect, the predicted trading in Table I differs from the actual trading of individual investors in two other ways. First, it involves partial adjustments to risky asset holdings, while, in reality, sales of entire positions are more common. Second, the share allocations require leverage, while, in reality, few individuals use leverage. The leverage is a consequence of our assumption that stock returns are binomially distributed. From the perspective of tractability, the binomial assumption is very useful, but, because it specifies a positive lower bound for the gross stock return, it leads to aggressive allocations. We have solved a two-period version of the decision problem in (10) to (12) using a *lognormal* return distribution and find that, in this case, the investor uses far less leverage. We discuss the lognormal case further in Section II.C.

Table I

An Example of Optimal Trading under Prospect Theory

We solve a portfolio problem with a risk-free asset and a binomially distributed risky asset. There are five dates, $t = 0, \ldots, 4$, and the interval between time 0 and time 4 is a year. The investor has prospect theory preferences defined over the trading profit he accumulates between time 0 and time 4. The top-left panel shows how the risky asset price evolves along a binomial tree. The top-right panel shows the state price density at each node in the tree. The bottom-left and bottom-right panels report, for each node, the optimal number of shares held in the risky asset and the optimal wealth, respectively. The investor's initial wealth is $40, the net risk-free rate is zero, and the initial price, annual net expected return, and annual standard deviation of the risky asset are $40, 0.1, and 0.3, respectively.

Risky Asset Price $P_{t,j}$					State Price Density $q_{t,j}$				
				72.9					0.46
			62.7					0.56	
		54.0		55.6			0.68		0.66
	46.5		47.9			0.83		0.80	
40		41.2		42.4	1		0.97		0.94
	35.5		36.5			1.18		1.14	
		31.4		32.4			1.38		1.34
			27.9					1.62	
				24.7					1.91
Risky Asset Shares Held $x_{t,j}$					Wealth $W_{t,j}$				
				–					163.39
			6.8					94.70	
		3.5		–			64.25		46.47
	1.8		0.5			50.75		42.87	
1.7		0.2		–	40		41.27		40.34
	1.5		0.0			32.45		40.15	
		2.7		–			26.26		40.02
			5.2					16.51	
				–					0

T-period stock price path for each of the investor's N_S stocks. We assume that all stocks have the same annual gross expected return μ and standard deviation σ, and that each one is distributed independently of the others. Given return process parameters, preference parameters, and the $10{,}000 \times N_S$ simulated stock price paths, we can use Proposition 1 to construct a data set of how the 10,000 prospect theory investors trade each of their N_S stocks over the T periods. For example, if one of an investor's stocks follows the

$$40 \rightarrow 46.5 \rightarrow 54.0 \rightarrow 47.9 \rightarrow 42.4$$

price path through the binomial tree in Table I, we know that the investor will hold 1.7, 1.8, 3.5, and 0.5 shares of the stock at trading dates $t = 0$, 1, 2, and 3, respectively.

To see if there is a disposition effect in our simulated data, we follow the method of Odean (1998) described in Section I. For each investor, we look at each of the $T - 1$ trading dates, $t = 1, \ldots, T - 1$. If the investor sells shares in any of his stocks at date $t \in \{1, \ldots, T-1\}$, we place each stock in his portfolio on that date into one of four categories. For every stock in his portfolio on date t that is *sold*, we count a realized gain if the stock price exceeds the average price at which the shares were purchased, and a realized loss otherwise. For every stock in the investor's portfolio on date t that is *not* sold, we count a paper gain if the stock price exceeds the average price at which the shares were purchased, and a paper loss otherwise. We count up the total number of paper gains and losses and realized gains and losses across all investors and all trading dates and compute the PGR and PLR ratios in equations (1) and (2). As in Odean (1998), we say that there is a disposition effect if PGR > PLR.

To implement this analysis, we fix the values of P_0, W_0, R_f, σ, α, λ, and N_S, and consider a range of values for μ and T. Specifically, we set the initial price of each stock to $P_0 = 40$, the initial wealth allocated to trading each stock by each investor to $W_0 = 40$, the gross risk-free rate to $R_f = 1$, the annual standard deviation of each stock to $\sigma = 0.3$, and the preference parameters for each investor to $(\alpha, \lambda) = (0.88, 2.25)$. Odean (1998) does not report the mean number of stocks held by the households in his sample, but Barber and Odean (2000), who use very similar data, report a mean value slightly above four. We therefore set $N_S = 4$. Our results are relatively insensitive to the value of N_S.

Table II reports PGR and PLR for various values of μ and T. Given a value for μ, a value for T, and the other parameter values from the previous paragraph,

Table II

Simulation Analysis of the Disposition Effect

For a given (μ, T) pair, we construct an artificial data set of how 10,000 investors trade stocks when they have prospect theory preferences defined over end-of-year stock-level trading profits; each investor trades four stocks, each stock has an annual gross expected return μ, and the year is divided into T trading periods. For each (μ, T) pair, we use the artificial data set to compute PGR and PLR, where PGR is the proportion of gains realized by all investors over the course of the year and PLR is the proportion of losses realized. The table reports "PGR/PLR" for each (μ, T) pair. An asterisk identifies a case in which there is no disposition effect (PGR < PLR). A dash indicates that the expected return μ is so low that the investor does not buy any stock at all.

Expected Return	Number of Trading Periods within the Year			
μ	$T = 2$	$T = 4$	$T = 6$	$T = 12$
1.03	–	–	–	0.55/0.51
1.04	–	–	0.52/0.55*	0.54/0.52
1.05	–	–	0.54/0.53	0.59/0.45
1.06	–	0.70/0.25	0.54/0.53	0.58/0.47
1.07	–	0.70/0.25	0.54/0.53	0.57/0.49
1.08	–	0.70/0.25	0.49/0.59*	0.47/0.60*
1.09	–	0.43/0.70*	0.49/0.59*	0.46/0.61*
1.10	0.0/1.0*	0.43/0.70*	0.49/0.59*	0.36/0.69*
1.11	0.0/1.0*	0.43/0.70*	0.49/0.59*	0.37/0.68*
1.12	0.0/1.0*	0.28/0.77*	0.24/0.81*	0.40/0.66*
1.13	0.0/1.0*	0.28/0.77*	0.24/0.83*	0.25/0.78*

we simulate an artificial data set and compute PGR and PLR for that data set. An asterisk identifies a case in which PGR is less than PLR, that is, a case in which the model fails to predict a disposition effect. Since the investors are loss-averse, they do not buy any stock at time 0 if the expected stock return is too low; these cases are indicated by dashes. The table shows that the threshold expected return at which investors buy the risky asset falls as the number of trading periods T rises. When there are many trading periods within the year, the kink in the utility function at time T is smoothed out from the perspective of time 0. This lowers investors' initial risk aversion and increases their willingness to buy the risky asset.

The table summarizes our analysis of the annual gain/loss implementation of prospect theory. The results are surprising. While this implementation does predict a disposition effect in some cases—in some cases, PGR does exceed PLR—we also see that, in many cases, PGR is *lower* than PLR. In other words, the annual gain/loss implementation of prospect theory often predicts the *opposite* of the disposition effect, namely that investors have a greater propensity to sell a stock trading at a paper *loss* than one trading at a paper gain.

For some readers, the most reasonable values of μ and T may be those that correspond to the top-right part of the table, where the disposition effect does hold. Even for these readers, however, there is an important conceptual point to take away from the table, a point that has not been noted in the literature to date: For *some* parameter values, the current implementation of prospect theory can predict the opposite of the disposition effect.

The table also shows us *when* the disposition effect is more likely to fail: when the expected risky asset return is high, and when the number of trading periods T is low. For example, when $T = 2$ the model always fails to predict a disposition effect, while for $T = 12$ it fails to do so in about half the cases we report. In the next section, we explain the intuition behind these findings.

B. An Example

To explain the results in Table II, we present a simple two-period example. In other words, we solve the problem in (10) to (12) for the case of $T = 2$, so that there are just three dates, $t = 0,\ 1$, and 2, and two allocation decisions, at $t = 0$ and $t = 1$. The two-period case is instructive because here, as Table II shows, the model always fails to predict a disposition effect, at least for the Tversky and Kahneman (1992) parameter values. As before, we set the gross risk-free rate to $R_f = 1$.

The two-period setting allows us to simplify the notation. We now use x_0 for the optimal share allocation at time 0, x_u for the optimal allocation at time 1 after a good stock return, and x_d for the optimal allocation at time 1 after a poor stock return. (The subscripts u and d refer to movements "up" or "down" the binomial tree.) The time 0 stock price is P_0 while $P_u = P_0 R_u$ and $P_d = P_0 R_d$ are the time 1 stock prices after a good stock return and after a poor stock return, respectively.

We will refer to the change in the investor's wealth between time 0 and time 1 as the time 1 gain/loss. It can take one of two values, ΔW_u or ΔW_d, depending on whether the stock goes up or down at time 1:

$$\Delta W_u = x_0 P_0(R_u - 1), \quad \Delta W_d = x_0 P_0(R_d - 1). \tag{30}$$

We will also be interested in the argument of the value function in equation (10). Since $T = 2$, we refer to this as the time 2 gain/loss. For $R_f = 1$, it equals the change in wealth between time 0 and time 2 and can take one of four values, ΔW_{uu}, ΔW_{ud}, ΔW_{du}, or ΔW_{dd}, depending on whether the stock goes up at time 1 and up at time 2, up then down, down then up, or down then down, respectively:

$$\begin{aligned} \Delta W_{uu} &= \Delta W_u + x_u P_u(R_u - 1) \qquad & \Delta W_{ud} &= \Delta W_u + x_u P_u(R_d - 1) \\ \Delta W_{du} &= \Delta W_d + x_d P_d(R_u - 1) \qquad & \Delta W_{dd} &= \Delta W_d + x_d P_d(R_d - 1). \end{aligned} \tag{31}$$

We say that there is a disposition effect in this two-period setting if and only if

$$x_u < x_0 \leq x_d. \tag{32}$$

In words, there is a disposition effect if the investor sells shares after a time 1 increase in the stock price ($x_u < x_0$) and buys shares or maintains the same position after a time 1 drop in the stock price ($x_0 \leq x_d$). It is straightforward to check that condition (32) is consistent with the definition of the disposition effect in Section II.A and in Odean (1998), namely that PGR > PLR.

In the example we now present, we set $(P_0, W_0) = (40, 40)$, $(\mu, \sigma) = (1.1, 0.3)$, and $(\alpha, \lambda) = (0.88, 2.25)$. When $T = 2$, this choice of μ and σ implies $(R_u, R_d) = (1.25, 0.85)$. Given these parameter values, we can use Proposition 1 to compute the investor's optimal trading strategy. We find that

$$(x_0, x_u, x_d) = (4.0, 5.05, 3.06).$$

Initially, then, the investor buys 4.0 shares of the risky asset. After a good stock return at time 1, he increases his position to 5.05 shares; and after a poor stock return at time 1, he decreases his position to 3.06 shares. Consistent with the $T = 2$ column of Table II, the investor's strategy is the opposite of the disposition effect. We now explain why.

Figure 2 plots the prospect theory value function in (6) for $(\alpha, \lambda) = (0.88, 2.25)$ and marks on the graph the time 1 and time 2 gains/losses defined in (30) and (31). This figure will be the focus of our discussion.

Since the time 0 allocation is $x_0 = 4.0$ shares, the investor's time 1 gain/loss after a good stock return is

$$\Delta W_u = x_0 P_0(R_u - 1) = (4.0)(40)(0.25) = 39.9.$$

This is point A. If the investor arrives at point A at time 1, we know that he increases his allocation to $x_u = 5.05$ shares. Points B and B' mark the time 2

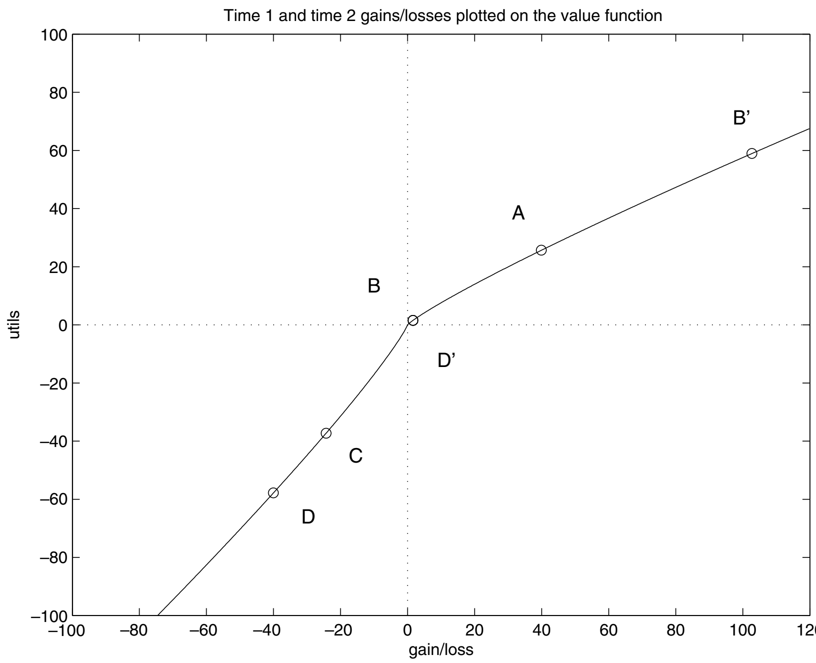

Figure 2. An example in which prospect theory fails to predict a disposition effect. The figure shows why an investor who derives prospect theory utility from the end-of-year profit he earns from trading a stock may exhibit the opposite of the disposition effect. There are three dates, $t = 0$, 1, and 2, and the interval between time 0 and time 2 is a year. The figure plots the Tversky and Kahneman (1992) prospect theory value function from Figure 1 and marks on it the various possible gains and losses in wealth at time 1 and time 2. If the stock does well at time 1, the investor moves to A. His optimal strategy is then to gamble to the edge of the concave region: If the stock does well (poorly) at time 2, he moves to point B' (B). If the stock instead does poorly at time 1, the investor moves to C. His optimal strategy is then to gamble to the edge of the convex region: If the stock does well (poorly) at time 2, he moves to point D' (D). Since the investor is loss-averse, the expected return on the stock needs to be reasonably high for him to buy it at all at time 0: A is therefore further from the vertical axis than C, and hence it takes a larger share allocation to gamble from A to the edge of the concave region than from C to the edge of the convex region. Thus, the investor is more likely to sell the stock after a loss, which is the opposite of the disposition effect.

gains/losses that this allocation could lead to. Specifically, point B' marks the time 2 gain/loss if the stock does well at time 2, namely

$$\Delta W_{uu} = \Delta W_u + x_u P_u (R_u - 1) = 39.9 + (5.05)(40)(1.25)(0.25) = 102.7,$$

while point B marks the time 2 gain/loss if the stock does poorly at time 2, namely

$$\Delta W_{ud} = \Delta W_u + x_u P_u (R_d - 1) = 39.9 + (5.05)(40)(1.25)(-0.15) = 1.6.$$

The figure shows that, after a gain at time 1, the investor takes a position in the stock such that, even if the stock subsequently does poorly, he still ends up with a slight gain at time 2 (point B). Put differently, after an initial gain, the investor gambles to the edge of the concave region.

Why does the investor follow this strategy? First, note that, since the investor is loss-averse, the expected return on the stock needs to be reasonably high for him to buy it at all at time 0. After a time 1 gain, he is in the concave region of the value function (point A). However, for the Tversky and Kahneman (1992) parameter values, the value function is only *mildly* concave over gains. Since the investor is almost risk-neutral in this region and since the expected stock return is reasonably high, he is willing to gamble almost to the edge of the concave region. He is not willing to take a larger gamble than this, however—in other words, a gamble that would bring point B to the left of the kink. If he did, he would risk ending up with a loss at time 2, which, given that he is loss-averse, would be very painful.

We now think about what happens if the stock does *poorly* at time 1. Given the time 0 allocation of $x_0 = 4.0$ shares, the investor's time 1 gain/loss after a poor stock return is

$$\Delta W_d = x_0 P_0(R_d - 1) = (4.0)(40)(-0.15) = -24.3.$$

This is point C. If the investor arrives at point C at time 1, we know that he decreases his allocation to $x_d = 3.06$ shares. Points D and D' mark the time 2 gains/losses that this allocation could lead to. Specifically, point D' marks the time 2 gain/loss if the stock does well at time 2, namely[6]

$$\Delta W_{du} = \Delta W_d + x_d P_d(R_u - 1) = -24.3 + (3.06)(40)(0.85)(0.25) = 1.6,$$

while point D marks the time 2 gain/loss if the stock does poorly at time 2, namely[7]

$$\Delta W_{dd} = \Delta W_d + x_d P_d(R_d - 1) = -24.3 + (3.06)(40)(0.85)(-0.15) = -40.$$

The figure shows that, after a loss at time 1, the investor takes a position in the stock such that, if the stock subsequently does well, he ends up with a small gain at time 2 (point D'). Why does he follow this strategy? After a time 1 loss, he is in the convex region of the value function (point C). He is therefore willing to gamble at least as far as the edge of the convex region. He is not willing to take a much larger gamble than this, however—in other words, a gamble that would bring point D' much to the right of the kink: To the right of the kink, the marginal utility of additional gains is significantly lower.

[6] Points B and D' are the same point. In our example, the optimal portfolio strategy is "path independence": The optimal wealth allocation at any date 2 node is independent of the path the stock price takes through the binomial tree to reach that node. The wealth allocation in the middle node at date 2 is therefore the same, whether the stock did well at date 1 and poorly at date 2, or vice versa.

[7] The time 2 gains/losses marked by D, B/D', and B' satisfy the prediction of Proposition 1 that the optimal final date wealth is either zero or a positive quantity that exceeds the reference wealth level. Here, the reference level is initial wealth, $W_0 = 40$. The \$40 loss at D therefore represents a final wealth of zero, while B/D' and B', by virtue of lying to the right of the kink, represent final wealth levels in excess of the reference level.

We can now complete the intuition for why, in this example, we obtain the opposite of the disposition effect. Since the investor is loss-averse, the stock must have a reasonably high expected return for him to buy it at all at time 0. In other words, $R_u - 1$ must be somewhat larger than $1 - R_d$. This means that the magnitude of the potential time 1 gain, $|x_0 P_0(R_u - 1)| = 39.9$, is *larger* than the magnitude of the potential time 1 loss, $|x_0 P_0(R_d - 1)| = 24.3$; in graphical terms, A is further from the vertical axis than C is. We know that after arriving at A, the investor gambles to the edge of the concave region. We also know that after arriving at C, he gambles to the edge of the convex region. However, since A is further from the vertical axis than C, it takes a larger share allocation to gamble from A to the edge of the concave region than it does to gamble from C to the edge of the convex region. The investor therefore takes more risk after a *gain* than after a loss. The propensity to sell is therefore *lower* after a gain than after a loss, contrary to the disposition effect.[8]

The fact that we obtain the opposite of the disposition effect is surprising because, at first sight, it seems that the annual gain/loss implementation of prospect theory *should* predict a disposition effect. The logic underlying this view—the flawed logic, as we soon explain—is this: Since the value function $v(\cdot)$ is concave over gains, an investor with a time 1 gain (point A) should take a relatively *small* gamble. Moreover, since the value function $v(\cdot)$ is convex over losses, an investor with a time 1 loss (point C) should gamble at least to the edge of the convex region, a relatively large gamble. It therefore seems that the investor should take less risk after a gain than after a loss, in other words, that he should have a greater propensity to sell a stock after a gain than after a loss and hence that the disposition effect should hold.

Where is the flaw in this argument? Since $v(\cdot)$ is only *mildly* concave in the region of gains, the only reason an investor would take a small position in the stock after a gain is if the expected stock return were unattractive; in other words, if it were only slightly higher than the risk-free rate. In this case, however, *the investor would not have bought the stock at time* 0! For him to buy the stock in the first place, its expected return must be reasonably high. But this, in combination with the mild concavity of $v(\cdot)$ in the region of gains, means that after a time 1 gain, the investor takes a *large* gamble, one that brings him almost to the edge of the concave region. Since this gamble is large, there is no disposition effect: The investor takes more risk after a gain than after a loss and therefore has a greater propensity to sell prior losers than prior winners.

This discussion also explains why, in Table II, the disposition effect *does* sometimes hold; specifically, when there are many trading periods T and the expected stock return μ is low. A key step in our explanation for why, in a two-period setting, the disposition effect fails is that, after a gain, the investor

[8] The intuition that the investor gambles to the edge of the concave region after a gain, or to the edge of the convex region after a loss, is appropriate when the stock return has a binomial distribution. We have also studied the case of a lognormally distributed stock return and find that, there, the investor uses strategies with a similar flavor: For example, after a gain at time 1, he takes a position in the stock so that much of the probability mass of the time 2 wealth distribution lies above the reference wealth level. We discuss the lognormal case in more detail in Section II.C.

gambles to the edge of the concave region. This relies on the fact that the expected stock return is quite high, which, in turn, is because otherwise, the investor would not buy the stock in the first place.

For large T, this logic can break down: When there are many trading periods before the final date, the kink in the time T utility function is smoothed over from the perspective of time 0, lowering the investor's initial risk aversion. He is therefore willing to buy the stock at time 0 even if its expected return is only *slightly* higher than the risk-free rate. When the expected return is this low, the concavity of $v(\cdot)$ in the region of gains leads the investor to take only a small position in the stock after a gain. As a result, the disposition effect can hold.

C. Robustness

Our analysis so far shows that, surprisingly, the annual gain/loss implementation of prospect theory often predicts the opposite of the disposition effect. How sensitive is this result to our modeling assumptions?

In our explanation for why our model can fail to predict a disposition effect, an important step was to note that, since the expected return on the stock must be high for the investor to buy it at all, the potential time 1 gain exceeds the potential time 1 loss in magnitude. This step is valid under our maintained assumption that $\pi = \frac{1}{2}$; in words, our assumption that, in each period, a good stock return and a poor stock return are equally likely. If $\pi > \frac{1}{2}$, however, a stock can have a high expected return by offering a small gain with high probability and a large loss with small probability. In this case, the disposition effect may hold: The share allocation needed to gamble to the edge of the concave region after the small gain is lower than the allocation needed to gamble to the edge of the convex region after the large loss. We caution, however, that individual stocks exhibit *positive* skewness in their returns rather than the negative skewness that this argument requires (Fama (1976)).

In our model, the risky asset return has a binomial distribution. This is a very useful assumption: For any number of periods T, it leads to an analytical solution for the optimal trading strategy and hence allows us to conduct the simulations summarized in Table II. At the same time, we want to be sure that our results are not special to the binomial case. We therefore solve the two-period version of the problem in (10) to (12) for the case where the risky asset return from time t to $t+1$, $R_{t,t+1}$, has a lognormal distribution, rather than a binomial one:

$$\log(R_{t,t+1} - \theta) = \mu + \sigma \varepsilon_{t,t+1}$$
$$\varepsilon_{t,t+1} \sim N(0,1), \quad \text{i.i.d. across periods,}$$

where θ, which satisfies $0 \le \theta < 1$, is the lowest gross return the risky asset can earn in any period.

Recall that, for the binomial distribution and the Tversky and Kahneman (1992) preference parameter values, the disposition effect never holds in the two-period case. We find that, for the lognormal distribution, the results are less extreme: There are some return process parameter values for which the disposition effect holds. However, once again, for a wide range of parameter values, there is no disposition effect. The intuition parallels that for the binomial case.

In our calculations, we have always set α and λ to the values estimated by Tversky and Kahneman (1992). What happens when we vary the values of α and λ? In the multi-period binomial model, but also in the two-period lognormal model, we find that, as the degree of loss aversion λ falls towards one, the disposition effect is less likely to hold. The reason is that, as λ falls, the investor takes a more aggressive, levered position in the risky asset at time 0. If the risky asset then does poorly, the investor cuts back on his holdings so as to prevent his final wealth from turning negative. This is the opposite of the disposition effect.

Our results also depend on α, which governs the curvature of the value function $v(\cdot)$, both in the region of gains and in the region of losses. When α falls substantially below the benchmark level of 0.88, we observe a disposition effect more often. It is easiest to see the intuition in the two-period binomial model. A lower α means greater concavity in the region of gains. This means that, after a time 1 gain, the investor takes a smaller position than before: He no longer gambles all the way to the edge of the concave region. Since a lower α also increases convexity over losses, the investor takes a larger position than before after a time 1 loss. Once α falls sufficiently—for $(\mu, \sigma) = (1.1, 0.3)$, once α falls below 0.77—the investor holds more shares after a loss than after a gain, and we obtain a disposition effect.

A very useful feature of our framework is that we can use the solution to the one-stock problem in (10) to (12) even in the multi-stock simulation of Section II.A. We can do this because we follow the prior literature in assuming narrow framing, so that the investor receives a separate component of utility from the trading profit on each of the stocks that he trades; and because we assume that the investor puts a limit, equal to W_0, on how much he is willing to lose between time 0 and time T from trading any one stock. An alternative assumption is that the investor instead puts a limit on how much he is willing to lose, *in total,* between time 0 and time T, from all of his stock trading activity. This alternative assumption significantly complicates the analysis, but we have been able to impose it in a two-stock, two-period model. We find that this model produces similar results: Once again, it often fails to predict a disposition effect.

The model of Section II assumes that, once the investor has decided, at time 0, on the maximum amount W_0 he is willing to lose from trading the risky asset, he sticks to that decision. We find that relaxing this assumption does not affect our conclusions. Specifically, suppose that, at time 1, the investor decides that he is willing to lose more than just the initial W_0 and that this decision is not anticipated at time 0. We find that, under this alternative assumption, the model again fails to predict a disposition effect.

Finally, in our model, the expected stock return is constant over time. If instead, after a time 1 gain, the investor for some reason lowers his estimate of the stock's expected return, he will be more inclined to sell and we may see a disposition effect after all.

The difficulty with this argument is that it requires that beliefs change in a way that cannot be considered rational: Odean (1998) finds that the average return of prior winners is *high*, not low, after they are sold. In this paper, we are investigating whether it is possible to understand the disposition effect without appealing to irrational beliefs. We therefore maintain a constant expected return throughout.

III. A Model That Applies Prospect Theory to Realized Gains and Losses

In the model of Section II, the investor receives prospect theory utility defined over annual stock-level trading profits. We now consider an alternative framework based on utility from *realized* gains and losses. In this case, if the investor buys shares of a stock at the start of the year and then, a few months later, sells some of the shares, he receives a jolt of prospect theory utility right then, at the moment of sale, where the argument of the prospect theory value function is the size of the realized gain or loss.

For simplicity, we work in a two-period model, so that there are three dates, $t = 0, 1$, and 2. As in Section II, there is a risk-free asset, which earns a gross return of $R_f = 1$ in each period, and a risky asset—a stock, say—whose gross return between time t and $t+1$, $R_{t,t+1}$, is given by (7) and (8) with $\pi = \frac{1}{2}$.

The difference between the model of this section and that of Section II lies in the investor's utility function. The investor now solves[9]

$$\max_{x_0,x_1} E_0\{v[(x_0 - x_1)(P_1 - P_0)]1_{\{x_1<x_0\}} + v[x_1(P_2 - P_b)]1_{\{x_1>0\}}\}, \tag{33}$$

where

$$P_b = \begin{cases} P_0 & \text{for} \quad x_1 \leq x_0 \\ \dfrac{x_0P_0 + (x_1 - x_0)P_1}{x_1} & \phantom{\text{for}} \quad x_1 > x_0 \end{cases}, \tag{34}$$

subject to

$$W_2 = W_0 + x_0P_0(R_{0,1} - 1) + x_1P_1(R_{1,2} - 1) \geq 0. \tag{35}$$

For $t \in \{0, 1, 2\}$, x_t, P_t, and W_t are the share allocation at time t, the risky asset price at time t, and wealth at time t, respectively. The variable P_b is the cost basis of any shares that the investor is still holding at time 2 and $1_{\{\}}$ is an indicator function that takes the value one if the condition in parentheses is satisfied and zero otherwise.

[9] For simplicity, our formulation assumes that the investor takes nonnegative positions in the risky asset. This will be true so long as the expected risky asset return exceeds the risk-free rate. It is straightforward to formulate the model in a way that accommodates short positions.

To understand the decision problem in (33) to (35), suppose that the investor buys x_0 shares at time 0. If, at time 1, he sells some shares—in other words, if $x_1 < x_0$—then he receives a jolt of prospect theory utility right then, at time 1. The argument of the value function $v(\cdot)$ is the size of the realized gain or loss, $(x_0 - x_1)(P_1 - P_0)$, or the number of shares sold multiplied by the difference between the sale price and the purchase price.

For simplicity, we assume that, at time 2, the investor sells any remaining shares in his possession. At that time, he receives prospect theory utility from the realized gain or loss, namely $x_1(P_2 - P_b)$, the number of shares sold multiplied by the difference between the sale price and the cost basis of the shares sold. Equation (34) says that, if the investor sells some shares at time 1, the cost basis of any remaining shares held until time 2 is still the initial purchase price, P_0. It also says that, if the investor buys additional shares at time 1, then the cost basis of the shares held until time 2 is the average price at which they were bought: A fraction x_0/x_1 were bought at a price of P_0 while a fraction $(x_1 - x_0)/x_1$ were bought at a price of P_1.

The time 1 state variables for the decision problem in (33) to (35) are x_0 and P_1. At time 1, then, the investor solves

$$J(x_0, P_1) = \max_{x_1 \in [0, W_1/(P_1(1-R_d))]} E_1\{v[(x_0 - x_1)(P_1 - P_0)]1_{\{x_1 < x_0\}} + v[x_1(P_2 - P_b)]1_{\{x_1 > 0\}}\}, \tag{36}$$

where $J(\cdot,\cdot)$ is the time 1 value function. To ensure that time 2 wealth is nonnegative, the time 1 share allocation can be at most $W_1/(P_1(1 - R_d))$. At time 0, the investor solves

$$\max_{x_0 \in [0, W_0/(P_0(1-R_d))]} E_0 J(x_0, P_1). \tag{37}$$

The time 0 share allocation can be at most $W_0/(P_0(1 - R_d))$ to ensure that time 1 wealth, and hence time 2 wealth, is nonnegative.

We solve (36) and (37) numerically. As in Section II, we set $(P_0, W_0) = (40, 40)$, $\sigma = 0.3$, $(\alpha, \lambda) = (0.88, 2.25)$, and consider several values of μ, the annual gross expected return on the risky asset. For given μ and σ, equations (28) and (29) with $T = 2$ allow us to compute R_u and R_d.

Panel B of Table III reports the optimal share holdings at time 0, the time 1 share holdings after a good risky asset return, and the time 1 share holdings after a poor risky asset return. In the notation of Section II.B, these three quantities are x_0, x_u, and x_d, respectively. A dash indicates a case in which the investor does not take a position in the risky asset at time 0. An asterisk indicates a case in which the model does not predict a disposition effect. For comparison, Panel A of Table III reports the results, using the same parameter values, for the two-period version of the annual gain/loss implementation of Section II.

The table shows that a model that applies prospect theory to realized gains and losses predicts a disposition effect more readily than the earlier model, which applied prospect theory to annual trading profit. At the same time, for

Table III

Optimal Share Allocations for Prospect Theory Investors

We solve a portfolio problem with a risk-free asset and a binomially distributed risky asset. There are three dates, $t = 0$, 1, and 2, and the interval between time 0 and time 2 is a year. Panel A corresponds to an investor who has prospect theory preferences defined over the trading profit he accumulates between time 0 and time 2. Panel B corresponds to an investor who has prospect theory preferences defined over realized gains and losses on the risky asset. The investor's initial wealth is \$40, the net risk-free rate is zero, and the initial price, annual gross expected return, and annual standard deviation of the risky asset are \$40, μ, and 0.3, respectively. x_0 is the optimal share allocation at time 0; x_u and x_d are the optimal time 1 share allocations after a good risky asset return and after a poor risky asset return, respectively. For example, when $\mu = 1.09$, an investor who derives prospect theory utility from realized gains and losses buys 3.4 shares of the risky asset at time 0. If the risky asset does well at time 1, he decreases his allocation to 2.6 shares. If the risky asset does poorly at time 1, he keeps the same allocation, namely 3.4 shares. An asterisk identifies a case in which there is no disposition effect. A dash indicates that the expected return on the risky asset is so low that the investor does not buy any of it at time 0.

Expected Return	Panel A			Panel B		
	Share Allocation			Share Allocation		
μ	x_0	x_u	x_d	x_0	x_u	x_d
1.06	–	–	–	–	–	–
1.07	–	–	–	–	–	–
1.08	–	–	–	–	–	–
1.09	–	–	–	3.4	2.6	3.4
1.10	4.0*	5.1*	3.1*	3.6	2.8	3.6
1.11	4.3*	5.7*	3.0*	3.7	3.0	3.7
1.12	4.6*	6.5*	3.0*	3.8*	5.5*	3.8*
1.13	4.9*	7.4*	3.0*	4.0*	6.0*	4.0*

high values of μ, we again observe the opposite of the disposition effect. Some of the intuition of the earlier model therefore carries over here as well.

Why does the realized gain/loss implementation of prospect theory predict a disposition effect more reliably than the annual gain/loss implementation? Note first that, for $R_f = 1$ and $T = 2$, the *sum* of the arguments of the two $v(\cdot)$ terms in (33) is equal in value to the argument of the $v(\cdot)$ term in equation (10), namely $W_2 - W_0$. The difference between the two models is therefore that, while the annual gain/loss model forces the investor to derive utility from the trading profit in a single lump at time 2, the realized gain/loss model allows the investor to split the trading profit into two components and to derive utility from each component separately, first at time 1 and then at time 2.

In the domain of losses, the investor in the realized gain/loss model would not want to divide the trading profit into two components: Since $v(\cdot)$ is convex over losses, losses are best experienced in one go, rather than in two separate pieces. In the domain of gains, however, the concavity of $v(\cdot)$ means that the investor is often keen to divide the trading profit into two pieces and to savor each one separately. After a gain at time 1, then, he often sells some shares.

IV. Related Research and Other Applications

The academic literature on the disposition effect starts with Shefrin and Statman (1985), who propose a framework with several elements: prospect theory, narrow framing/mental accounting, utility from realized gains and losses, regret utility, and a self-control problem. Each element is designed to explain a separate piece of empirical evidence: Prospect theory defined over realized gains and losses explains the basic disposition effect; narrow framing explains why investors don't like "tax swaps," in other words, selling their position in a losing stock and immediately transferring the proceeds to another, similar stock; regret utility explains why some individual investors do not display a disposition effect; and the self-control problem explains some rules of thumb used by professional traders.

Our analysis provides new support for Shefrin and Statman's (1985) decision to implement prospect theory over *realized* gains and losses. The assumption that the investor derives utility from realized gains and losses is a significant departure from the traditional framework, but without it, it is much harder to generate a disposition effect: The analysis in Section II shows that a model that applies prospect theory to *annual* gains and losses does not predict a disposition effect very reliably.

Hens and Vlcek (2005) also investigate the link between prospect theory and the disposition effect, albeit in a somewhat different framework. They consider a two-period model with three dates, $t = 0,\ 1$, and 2, and two assets, a risk-free asset and a risky asset. At time 1, the investor receives prospect theory utility defined over the trading profit earned between time 0 and time 1. At time 2, he receives prospect theory utility defined over the total trading profit earned between time 0 and time 2. Hens and Vlcek (2005) assume that the investor acts myopically: At time 0, he chooses a share allocation to maximize time 1 utility; at time 1, he chooses a share allocation to maximize time 2 utility. Echoing our own results, the authors find that this model often fails to predict a disposition effect.

The analysis in Hens and Vlcek (2005) complements the analysis in our paper: We consider some dimensions that they do not and they consider some dimensions that we do not. For example, in our analysis of the annual gain/loss implementation of prospect theory, we allow for any number of trading periods, rather than just two, and for full dynamic optimization, rather than just myopic decision-making. Our ability to explore cases with many trading periods turns out to be useful: Qualitatively, the results are different for high T in that the disposition effect tends to hold more often. We also study a realized gain/loss implementation of prospect theory, which Hens and Vlcek (2005) do not.

At the same time, Hens and Vlcek's (2005) simplifying assumption of myopic decision-making allows them to take prospect theory's probability weighting feature into account. They find that probability weighting plays only a minor role in determining whether prospect theory predicts a disposition effect. It plays a larger role in determining whether prospect theory predicts a related concept, the "ex post disposition effect"—the investor's propensity to realize gains as opposed to losses, *without* conditioning on the initial purchase decision.

Gomes (2005) studies a two-period economy in which some investors have preferences that are related to, but different from, prospect theory. Specifically, for losses below some specific point, the convex section of the prospect theory value function is replaced with a concave segment. Gomes (2005) is primarily interested in volume and volatility but also includes a short discussion of the disposition effect. He finds that, for a particular range of preference parameter values, the investors in his framework do exhibit a disposition effect. Our analysis shows that this result may be special to his model: For unmodified prospect theory and for the Tversky and Kahneman (1992) parameter values, a two-period model that defines utility over annual gains and losses always fails to predict a disposition effect. As with Hens and Vlcek (2005), Gomes (2005) does not explore beyond the two-period case.

Kyle, Ou-Yang, and Xiong (2006) consider an investor who is endowed with a project, or indivisible asset, and who is trying to decide when to liquidate the project. On liquidation, the investor receives prospect theory utility defined over the difference between the project's liquidation value and the amount invested in the project. This analysis differs from ours in a number of ways. Most importantly, it does not take the investor's initial buying decision into account. As soon as we do, we recognize that the expected risky asset return must exceed a certain level. This, in turn, affects the likelihood that prospect theory will predict a disposition effect.

So far, we have applied our analysis in one particular context: the trading of individual stocks. However, researchers have also uncovered disposition effects in other settings. Genesove and Mayer (2001) find that homeowners are reluctant to sell their houses at prices below the original purchase price. Heath, Huddart, and Lang (1999) find that executives are more likely to exercise their stock options when the underlying stock price exceeds a reference point, the stock's highest price over the previous year, than when it falls below that reference point. Coval and Shumway (2005) show that futures traders who have accumulated trading profits by the midpoint of the day take *less* risk in the afternoon than traders who, by the midpoint of the day, have trading losses.

Our analysis can be applied in all of these settings. A model that defines prospect theory over the gains and losses an investor earns over a *fixed* interval, whether a day or a year, often predicts the opposite of the disposition effect. Such a model is therefore not well suited to explaining the above findings. A model that defines prospect theory over realized gains and losses explains the evidence more readily.

V. Conclusion

In this paper, we investigate whether prospect theory preferences can predict a disposition effect. We consider two implementations of prospect theory: In one case, preferences are defined over *annual* gains and losses; in the other, they are defined over *realized* gains and losses. Surprisingly, the annual gain/loss model often fails to predict a disposition effect. The realized gain/loss model, however, predicts a disposition effect more reliably.

The idea that investors might derive utility from the act of realizing a gain or loss on a specific asset that they own has not received much attention in the finance literature to date. Our analysis shows that, while an unusual feature of preferences, utility from realized gains and losses nonetheless offers a simple way of thinking about a puzzling phenomenon, the disposition effect. This suggests that it may be useful to conduct a more comprehensive analysis of realized gain/loss utility and to see whether it can shed light on other aspects of investor trading, or even on asset prices. Barberis and Xiong (2008) take a first step in this direction.

Appendix

Proof of Proposition 1: We use the insight of Cox and Huang (1989) that, when markets are complete, an investor's dynamic optimization problem can be rewritten as a *static* problem in which the investor directly allocates wealth across final period states.

When the investor's utility function is concave, the final period wealth allocation is not path dependent: The optimal wealth allocation to node j at time T does not depend on the path the stock price takes through the binomial tree before arriving at that node. In our case, however, the investor has a prospect theory utility function, which is not concave. His final period wealth allocation could therefore be path dependent. To accommodate this possibility, we allow the investor to allocate wealth across stock price *paths*. Later, we will argue that it is reasonable to restrict our attention to final period wealth allocations that are *not* path dependent.

There are $M = 2^T$ paths that the stock price can take to reach one of the date T nodes. We denote these paths by $i \in \{1, 2, \ldots, M\}$. The ex-ante probability of path i is $\hat{\pi}_i$, so that

$$\sum_{i=1}^{M} \hat{\pi}_i = 1.$$

The price of a contingent claim that pays \$1 at time T if the stock price evolves along path i is $\hat{p}_i$. The state price density at the endpoint of the path is therefore $\hat{q}_i = \hat{p}_i / \hat{\pi}_i$. In addition, W_0 is the investor's initial wealth at time 0, R_f is the per-period gross risk-free rate, and $\{\hat{W}_i\}_{i=1}^{M}$ are the investor's wealth allocations at the end of each path. Hats indicate variables that are indexed by path, rather than by node. While the optimal date T wealth allocations may be path dependent, the state price densities are not: If paths i and j end at the same date T node, then $\hat{q}_i = \hat{q}_j$. We compute the state price density explicitly later in the proof.

Applying the reasoning of Cox and Huang (1989), we can rewrite the problem in (10) to (12) as

$$V = \max_{\{\hat{W}_i\}} \sum_{i=1}^{M} \hat{\pi}_i v\big(\hat{W}_i - W_0 R_f^T\big), \tag{A1}$$

subject to the budget constraint

$$\sum_{i=1}^{M} \hat{\pi}_i \hat{q}_i \hat{W}_i = W_0 \tag{A2}$$

and the nonnegativity of wealth constraint

$$\hat{W}_i \geq 0, \quad 1 \leq i \leq M. \tag{A3}$$

We write the reference level of wealth $W_0 R_f^T$ as $\bar{W}$, for short, and define $\hat{w}_i = \hat{W}_i - \bar{W}$ to be the investor's gain/loss relative to that reference level. The problem in (A1) to (A3) then becomes[10]

$$V = \max_{\{\hat{w}_i\}} \sum_{i=1}^{M} \hat{\pi}_i v(\hat{w}_i), \tag{A4}$$

subject to

$$\sum_{i=1}^{M} \hat{\pi}_i \hat{q}_i \hat{w}_i = 0 \tag{A5}$$

$$\hat{w}_i \geq -\bar{W}, \quad 1 \leq i \leq M. \tag{A6}$$

We now prove the proposition through a series of lemmas.

LEMMA A1: *There exists at least one optimum.*

Proof of Lemma A1: The set of feasible $\{\hat{w}_i\}$ defined by constraints (A5) and (A6) is compact. The existence result then follows directly from Weierstrass's theorem. Q.E.D.

We now describe some of the properties of the optimum. Without loss of generality, we assume

$$\hat{\pi}_1^{1-\alpha} \hat{q}_1^{-\alpha} \leq \hat{\pi}_2^{1-\alpha} \hat{q}_2^{-\alpha} \leq \cdots \leq \hat{\pi}_M^{1-\alpha} \hat{q}_M^{-\alpha}.$$

LEMMA A2: *If* $\hat{\pi}_M^{1-\alpha} \hat{q}_M^{-\alpha} > \lambda \hat{\pi}_1^{1-\alpha} \hat{q}_1^{-\alpha}$, $\{\hat{w}_i = 0\}_{i=1}^{M}$ *cannot be the optimum.*

[10] In what follows, the term "wealth allocation" sometimes refers to a wealth level, such as $\hat{W}_i$, and sometimes to a gain/loss, such as $\hat{w}_i$. It will always be clear from the context which of the two meanings applies.

Proof of Lemma A2: We prove the lemma by contradiction. Suppose that $\{\hat{w}_i = 0\}_{i=1}^{M}$ is the optimum, so that the investor's value function takes the value $V = 0$. Consider the strategy

$$\hat{w}_1 = -x, \hat{w}_2 = \cdots = \hat{w}_{M-1} = 0, \hat{w}_M = \frac{\hat{\pi}_1 \hat{q}_1}{\hat{\pi}_M \hat{q}_M} x,$$

where $x \in (0, \bar{W}]$. By construction, this strategy satisfies the budget constraint. The associated value function is

$$V' = \hat{\pi}_M \frac{\hat{\pi}_1^{\alpha} \hat{q}_1^{\alpha}}{\hat{\pi}_M^{\alpha} \hat{q}_M^{\alpha}} x^{\alpha} - \lambda \hat{\pi}_1 x^{\alpha} = \left(\frac{\hat{\pi}_M^{1-\alpha} \hat{q}_M^{-\alpha}}{\hat{\pi}_1^{1-\alpha} \hat{q}_1^{-\alpha}} - \lambda \right) \hat{\pi}_1 x^{\alpha} > V = 0.$$

Thus, we obtain a contradiction; $\{\hat{w}_i = 0\}_{i=1}^{M}$ cannot be the optimum. Q.E.D.

LEMMA A3*: If the investor's optimal gain/loss $\hat{w}_i$ is different from zero at the end of some path, then it is different from zero at the end of all paths.*

Proof of Lemma A3: We prove the lemma by contradiction. Suppose that the investor's gain/loss is zero at the end of path i, so that $\hat{w}_i = 0$. If the investor's gain/loss is negative at the end of one path, it must be positive at the end of another path. We therefore assume, without loss of generality, that the investor's gain/loss is positive at the end of path j, so that $\hat{w}_j = x > 0$. The contribution of paths i and j to total utility is $J = \hat{\pi}_j x^{\alpha}$. We now modify this strategy by moving a small amount of wealth $\delta > 0$ from path j to path i, so that

$$\hat{w}_i = \delta, \ \hat{w}_j = x - \frac{\hat{\pi}_i \hat{q}_i}{\hat{\pi}_j \hat{q}_j} \delta.$$

The contribution of paths i and j to total utility is now

$$J(\delta) = \hat{\pi}_i \delta^{\alpha} + \hat{\pi}_j \left(x - \frac{\hat{\pi}_i \hat{q}_i}{\hat{\pi}_j \hat{q}_j} \delta \right)^{\alpha}.$$

It is straightforward to verify that

$$J'(0) > 0,$$

so that moving wealth from path j to path i increases the investor's value function. This contradicts the initial assumption that $\hat{w}_i = 0$ is optimal. Hence, the optimal gain/loss is different from zero at the end of all paths. Q.E.D.

LEMMA A4*: If the optimal allocation is nonzero, there must be at least one path at the end of which the gain/loss is $-\bar{W}$, so that the investor is wealth constrained.*

Proof of Lemma A4: Suppose that $\vec{w} = (\hat{w}_1, \hat{w}_2, \ldots, \hat{w}_M)$ is a nonzero optimal allocation, so that the value function $J(\vec{w}) > J(0) = 0$. Suppose also that the wealth constraint is never binding, so that $\hat{w}_i > -\bar{W}, \ \forall i$. This implies that there exists $k > 1$ such that $k\vec{w}$ is a feasible allocation, which, in turn, means that $J(k\vec{w}) = k^{\alpha} J(\vec{w}) > J(\vec{w})$. Thus, we have a contradiction. Q.E.D.

In any nontrivial optimum—any optimum in which $\hat{w}_i$ does not equal zero for all i—there are three possible wealth allocations at the end of a path: a positive allocation ($\hat{w}_i > 0$), an unconstrained negative allocation ($-\bar{W} < \hat{w}_i < 0$), or a constrained negative allocation ($\hat{w}_i = -\bar{W}$). In particular, from Lemma A3, we know that $\hat{w}_i = 0$ cannot be an optimal allocation.

To solve for the optimal allocation, we use the Lagrange multiplier method. The Lagrangian is

$$L = \sum_{i=1}^{M} \hat{\pi}_i v(\hat{w}_i) - \mu_0 \sum_{i=1}^{M} \hat{\pi}_i \hat{q}_i \hat{w}_i + \sum_{i=1}^{M} \mu_i (\hat{w}_i + \bar{W}),$$

where $\mu_0 > 0$ is the multiplier associated with the budget constraint and $\mu_i \geq 0$ is the multiplier associated with the wealth constraint on path i. The first-order condition for $\hat{w}_i$ is

$$v'(\hat{w}_i) = \mu_0 \hat{q}_i - \mu_i / \hat{\pi}_i, \quad 1 \leq i \leq M.$$

Since μ_i is associated with an inequality constraint,

$$\mu_i = 0 \text{ if } \hat{w}_i > -\bar{W}, \quad 1 \leq i \leq M,$$

and

$$\mu_i > 0 \text{ if } \hat{w}_i = -\bar{W}, \quad 1 \leq i \leq M.$$

Since $v'(\cdot)$ ranges from zero to infinity in both the negative and positive domains, there are three possible solutions to the first-order condition:

1. A positive wealth allocation:

$$\hat{w}_i = \left(\frac{\alpha}{\mu_0 \hat{q}_i} \right)^{\frac{1}{1-\alpha}}. \tag{A7}$$

2. An unconstrained negative wealth allocation:

$$\hat{w}_i = -\left(\frac{\alpha\lambda}{\mu_0 \hat{q}_i} \right)^{\frac{1}{1-\alpha}}, \text{ if } \hat{q}_i > \frac{\alpha\lambda}{\mu_0} \bar{W}^{-(1-\alpha)}. \tag{A8}$$

3. A constrained negative wealth allocation:

$$\hat{w}_i = -\bar{W}, \text{ if } \hat{q}_i > \frac{\alpha\lambda}{\mu_0} \bar{W}^{-(1-\alpha)}. \tag{A9}$$

For any solution to the first-order condition, we sort the M paths based on their respective wealth allocations: paths $\{1, \ldots, k\}$ have a positive allocation, paths $\{k+1, \ldots, m\}$ have an unconstrained negative allocation, and paths $\{m+1, \ldots, M\}$ have a constrained negative allocation.

The multiplier μ_0 can be determined from the budget constraint

$$\sum_{i=1}^{k} \hat{\pi}_i \hat{q}_i \hat{w}_i + \sum_{j=k+1}^{m} \hat{\pi}_j \hat{q}_j \hat{w}_j = \sum_{l=m+1}^{M} \hat{\pi}_l \hat{q}_l \bar{W}.$$

Substituting in the values of $\hat{w}_i$, we obtain

$$\left(\frac{\alpha}{\mu_0}\right)^{\frac{1}{1-\alpha}} = \frac{\bar{W} \sum_{l=m+1}^{M} \hat{\pi}_l \hat{q}_l}{\sum_{i=1}^{k} \hat{\pi}_i \hat{q}_i^{-\frac{\alpha}{1-\alpha}} - \lambda^{\frac{1}{1-\alpha}} \sum_{j=k+1}^{m} \hat{\pi}_j \hat{q}_j^{-\frac{\alpha}{1-\alpha}}}. \tag{A10}$$

The value function is therefore

$$\begin{aligned} V &= \sum_{i=1}^{k} \hat{\pi}_i \hat{w}_i^{\alpha} - \sum_{j=k+1}^{m} \lambda \hat{\pi}_j (-\hat{w}_j)^{\alpha} - \sum_{l=m+1}^{M} \lambda \hat{\pi}_l \bar{W}^{\alpha} \\ &= \bar{W}^{\alpha} \left[\left(\sum_{l=m+1}^{M} \hat{\pi}_l \hat{q}_l \right)^{\alpha} \left(\sum_{i=1}^{k} \hat{\pi}_i \hat{q}_i^{-\frac{\alpha}{1-\alpha}} - \lambda^{\frac{1}{1-\alpha}} \sum_{j=k+1}^{m} \hat{\pi}_j \hat{q}_j^{-\frac{\alpha}{1-\alpha}} \right)^{1-\alpha} - \lambda \sum_{l=m+1}^{M} \hat{\pi}_l \right]. \end{aligned} \tag{A11}$$

The optimal date T wealth allocation is determined by comparing all possible solutions. Such a comparison reveals several additional properties of the optimal allocation.

LEMMA A5: *It is not optimal to have a path with an unconstrained negative allocation,* $\hat{w}_i : -\bar{W} < \hat{w}_i < 0$.

Proof of Lemma A5: From equation (A11), we see that replacing an unconstrained negative wealth allocation with a positive wealth allocation strictly improves the value function. Thus, it is not optimal to have a path with an unconstrained negative allocation. Q.E.D.

Lemma A5 implies that we can write the value function as

$$V = \bar{W}^{\alpha} \left[\left(\sum_{i=1}^{k} \hat{\pi}_i \hat{q}_i^{-\frac{\alpha}{1-\alpha}} \right)^{1-\alpha} \left(\sum_{l=k+1}^{M} \hat{\pi}_l \hat{q}_l \right)^{\alpha} - \lambda \sum_{l=k+1}^{M} \hat{\pi}_l \right]. \tag{A12}$$

LEMMA A6: *Suppose that, in each period, a good stock return and a poor stock return are equally likely. Then any path with a constrained negative wealth allocation must have a state price density not lower than that of any path with a positive allocation.*

Proof of Lemma A6: Given the equal probability of a good or poor return in each period, each price path has the same probability: $\hat{\pi}_i = 2^{-T},\ i = 1, \ldots, M$. The value function is then

$$V = 2^{-T}\bar{W}^{\alpha}\left[\left(\sum_{i=1}^{k}\hat{q}_i^{-\frac{\alpha}{1-\alpha}}\right)^{1-\alpha}\left(\sum_{l=k+1}^{M}\hat{q}_l\right)^{\alpha} - \lambda(M-k)\right]. \qquad \text{(A13)}$$

Suppose that $\hat{q}_l < \hat{q}_i$ for some $i \in \{1, \ldots, k\}$ and $l \in \{k+1, \ldots, M\}$. Equation (A13) shows that assigning path l a positive wealth allocation and path i a constrained negative allocation $-\bar{W}$ strictly improves the value function. Thus, we obtain a contradiction. Any path with a constrained negative allocation must therefore have a state price density not lower than that of any path with a positive allocation. Q.E.D.

Lemma A6 shows that the optimal wealth allocation has a threshold property: Paths with a state price density higher than a certain level have a constrained negative allocation, while paths with a state price density lower than that level have a positive allocation. This threshold property may not hold if the probabilities of a good or poor return in each period are not the same.

Another corollary of Lemma A6 is that the final date wealth allocations can be path dependent at, at most, one node. If they were path dependent at more than one node, we could find a path with a constrained negative wealth allocation that had a state price density lower than that of a path with a positive wealth allocation. This would contradict the lemma.

In Barberis and Xiong (2006), the NBER working paper version of this paper, we show that if the wealth allocation is path dependent at even one final date node, it is impossible to clear markets in a simple equilibrium model with both expected utility investors and prospect theory investors. Since, in our view, this is an undesirable property, we now restrict our attention to path independent wealth allocations. We can therefore revert to the notation of Section II where $P_{t,i}$, $W_{t,i}$, $x_{t,i}$, and $q_{t,i}$ denote the stock price, optimal wealth allocation, optimal share allocation, and state price density in node i at date t.

The final ingredient we need to complete the proof is the state price density $q_{t,i}$. Since the price process for the risky asset is homogeneous, the state price process $q_{t,i}$ must also be homogeneous. We therefore assume that, from one period to the next, $q_{t,i}$ changes either by a factor of q_u or by a factor of q_d, where, as usual, the subscripts u and d refer to movements up and down the binomial tree, respectively. A standard property of the state price density is that it correctly prices the stock and the risk-free asset at each node, given next-period prices:

$$P_{t,i} = \frac{\frac{1}{2}q_{t+1,i}P_{t+1,i} + \frac{1}{2}q_{t+1,i+1}P_{t+1,i+1}}{q_{t,i}}$$

$$1 = \frac{\frac{1}{2}q_{t+1,i}R_f + \frac{1}{2}q_{t+1,i+1}R_f}{q_{t,i}},$$

which are equivalent to

$$2 = q_u R_u + q_d R_d$$
$$\frac{2}{R_f} = q_u + q_d.$$

We therefore obtain

$$q_u = \frac{2(R_f - R_d)}{R_f(R_u - R_d)}, \qquad q_d = \frac{2(R_u - R_f)}{R_f(R_u - R_d)}. \tag{A14}$$

Since $(R_f - R_d) < (R_u - R_f)$, we have $q_u < q_d$.

We can now complete the proof. From (A14), we know that the state price density increases as we go down the $T+1$ date T nodes, so that $q_{T,1} < q_{T,2} < \cdots < q_{T,T+1}$. From Lemma A6, equation (A7), and equation (A10), and remembering that we are now summing over nodes, not paths, we know that, for the k^* top nodes in the final period, where $1 \leq k^* \leq T$, the investor chooses an optimal wealth level of

$$W_{T,i} = \bar{W}\left[1 + \frac{q_{T,i}^{-\frac{1}{1-\alpha}} \sum_{l=k^*+1}^{T+1} q_{T,l}\pi_{T,l}}{\sum_{l=1}^{k^*} q_{T,l}^{-\frac{\alpha}{1-\alpha}} \pi_{T,l}}\right], \quad i \leq k^*,$$

where $\pi_{T,l}$ is the probability of reaching node l on date T:

$$\pi_{T,l} = \frac{T!2^{-T}}{(T-l+1)!(l-1)!}.$$

For the bottom $T+1-k^*$ nodes, he chooses an optimal wealth level of zero.

To determine k^*, we compute the investor's utility for each of the T possible values of k^*, $1 \leq k \leq T$, and look for the wealth allocation strategy that maximizes utility. Suppose that the investor chooses a positive wealth allocation for the top k nodes. From equation (A12), utility is then given by

$$V = \bar{W}^\alpha\left[\left(\sum_{l=1}^{k} q_{T,l}^{-\frac{\alpha}{1-\alpha}}\pi_{T,l}\right)^{1-\alpha}\left(\sum_{l=k+1}^{T+1} q_{T,l}\pi_{T,l}\right)^{\alpha} - \lambda\sum_{l=k+1}^{T+1}\pi_{T,l}\right];$$

k^* is the value of k that maximizes this utility.

Given the optimal wealth levels in the final period, $W_{T,j}$, we can compute optimal wealth levels at all earlier dates using the state price density, as shown in equation (26).

The final step is to compute optimal share holdings at each node. Suppose that, in node i at date t, the investor holds $x_{t,i}$ shares of stock and B dollars of the risk-free asset. His wealth in node i at date $t+1$ will therefore be

$x_{t,i}P_{t+1,i} + BR_f$, and in node $i+1$ at date $t+1$, $x_{t,i}P_{t+1,i+1} + BR_f$. The difference must equal $W_{t+1,i} - W_{t+1,i+1}$, so that

$$x_{t,i} = \frac{W_{t+1,i} - W_{t+1,i+1}}{P_{t+1,i} - P_{t+1,i+1}} = \frac{W_{t+1,i} - W_{t+1,i+1}}{P_0\left(R_u^{t-i+2}R_d^{i-1} - R_u^{t-i+1}R_d^{i}\right)},$$

which is equation (25). This completes the proof of Proposition 1. Q.E.D.

REFERENCES

Barber, Brad, and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock performance of individual investors, *Journal of Finance* 55, 773–806.

Barberis, Nicholas, and Ming Huang, 2008, Stocks as lotteries: The implications of probability weighting for security prices, *American Economic Review* 98, 2066–2100.

Barberis, Nicholas, and Wei Xiong, 2006, What drives the disposition effect? An analysis of a long-standing preference-based explanation, NBER Working paper 12397.

Barberis, Nicholas, and Wei Xiong, 2008, Realization utility, Working paper, Yale University.

Benartzi, Shlomo, and Richard Thaler, 1995, Myopic loss aversion and the equity premium puzzle, *Quarterly Journal of Economics* 110, 75–92.

Berkelaar, Arjan, Roy Kouwenberg, and Thierry Post, 2004, Optimal portfolio choice under loss aversion, *Review of Economics and Statistics* 86, 973–987.

Coval, Joshua, and Tyler Shumway, 2005, Do behavioral biases affect prices? *Journal of Finance* 60, 1–34.

Cox, John, and Chi-fu Huang, 1989, Optimal consumption and portfolio policies when asset prices follow a diffusion process, *Journal of Economic Theory* 39, 33–83.

Dhar, Ravi, and Ning Zhu, 2006, Up close and personal: An individual level analysis of the disposition effect, *Management Science* 52, 726–740.

Fama, Eugene, 1976, *Foundations of Finance* (Basic Books, New York).

Feng, Lei, and Mark Seasholes, 2005, Do investor sophistication and trading experience eliminate behavioral biases in finance markets? *Review of Finance* 9, 305–351.

Frazzini, Andrea, 2006, The disposition effect and underreaction to news, *Journal of Finance* 61, 2017–2046.

Genesove, David, and Christopher Mayer, 2001, Loss aversion and seller behavior: Evidence from the housing market, *Quarterly Journal of Economics* 116, 1233–1260.

Gomes, Francisco, 2005, Portfolio choice and trading volume with loss averse investors, *Journal of Business* 78, 675–706.

Grinblatt, Mark, and Bing Han, 2005, Prospect theory, mental accounting, and the disposition effect, *Journal of Financial Economics* 78, 311–339.

Grinblatt, Mark, and Matti Keloharju, 2001, What makes investors trade? *Journal of Finance* 56, 589–616.

Heath, Chip, Steven Huddart, and Mark Lang, 1999, Psychological factors and stock option exercise, *Quarterly Journal of Economics* 114, 601–627.

Hens, Thorsten, and Martin Vlcek, 2005, Does prospect theory explain the disposition effect? Working paper, University of Zurich.

Kahneman, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47, 263–291.

Kyle, Albert, Hui Ou-Yang, and Wei Xiong, 2006, Prospect theory and liquidation decisions, *Journal of Economic Theory* 129, 273–285.

Odean, Terrance, 1998, Are investors reluctant to realize their losses? *Journal of Finance* 53, 1775–1798.

Shefrin, Hersh, and Meir Statman, 1985, The disposition to sell winners too early and ride losers too long, *Journal of Finance* 40, 777–790.

Tversky, Amos, and Daniel Kahneman, 1992, Advances in prospect theory: Cumulative representation of uncertainty, *Journal of Risk and Uncertainty* 5, 297–323.