WHAT DRIVES THE DISPOSITION EFFECT? AN ANALYSIS OF A LONG-STANDING PREFERENCE-BASED EXPLANATION

Nicholas Barberis
Wei Xiong

Working Paper **12397**

What Drives the Disposition Effect? An Analysis of a Long-Standing Preference-Based
Explanation
Nicholas Barberis and Wei Xiong
NBER Working Paper No. 12397
July 2006
JEL No. G11, G12

ABSTRACT

One of the most striking portfolio puzzles is the "disposition effect": the tendency of individuals to
sell stocks in their portfolios that have risen in value since purchase, rather than fallen in value.
Perhaps the most prominent explanation for this puzzle is based on prospect theory. Despite its
prominence, this explanation has received little formal scrutiny. We take up this task, and analyze
the trading behavior of investors with prospect theory preferences. We find that, at least for the
simplest implementation of prospect theory, the link between these preferences and the disposition
effect is not as obvious as previously thought: in some cases, prospect theory does indeed predict a
disposition effect, but in others, it predicts the opposite. We provide intuition for these results, and
identify the conditions under which the disposition effect holds or fails. We also discuss the
implications of our results for other disposition-type effects that have been documented in settings
such as the housing market, futures trading, and executive stock options.

Nicholas C. Barberis
Yale School of Management
135 Prospect Street PO Box 208200
New Haven CT 06520-8200
and NBER
nick.barberis@yale.edu

Wei Xiong
Princeton University
Bendheim Center for Finance
26 Prospect Avenue
Princeton, NJ 08450
and NBER
wxiong@princeton.edu

1   Introduction

One of the most robust facts about the trading of individual investors is the so-called "disposition effect": the finding that, when an individual investor sells a stock in her portfolio, she has a greater propensity to sell a stock that has gone *up* in value since purchase, than one that has gone down.

Robust though the disposition effect is, its cause remains unclear. *Why* do individual investors prefer to 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 many of the most obvious potential explanations are inconsistent with the evidence. Take, for example, the most obvious hypothesis of all, the "information" hypothesis: individual investors sell stocks with paper gains because they have private information that these stocks will subsequently do poorly; and they hold on to stocks with paper losses because they think that these stocks will rebound. The problem with this story, as Odean (1998) points out, is that the subsequent return of the prior winners people sell, is, on average, *higher* than the subsequent return of the prior losers they hold on to. Odean (1998) goes on to show that other standard hypotheses based on taxes, rebalancing, or transaction costs also fail to capture important features of the data.

In light of the difficulties faced by these standard theories, an alternative hypothesis based on prospect theory has gained favor. Prospect theory, a prominent theory of decision-making under risk proposed by Kahneman and Tversky (1979) and refined in 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 in the region of gains and convex in the region of losses, like that in Figure 1. This functional form captures the experimental finding that people, on the one hand, are risk averse over gains – they prefer a certain $100 to a 50:50 bet to win $0 or $200 – but, on the other hand, are risk-*seeking* over losses: they prefer a 50:50 bet to lose $0 or $200 to a certain loss of $100. The value function is also kinked at the origin, a feature known as loss aversion. This captures a greater sensitivity to losses – even small losses – than to gains of the same magnitude: a 50:50 bet to win $110 or lose $100, for example, is typically rejected.

Prospect theory appears to offer a simple way of understanding the disposition effect. If an individual investor is risk averse over gains, she should be inclined to sell a stock that is trading at a gain, in other words, a stock that has risen since purchase; and if she is risk-seeking over losses, she should be inclined to hold on to a stock that is trading at a loss. Researchers have been linking prospect theory and the disposition effect in this way for over 20 years. A review of the literature suggests that the prospect theory view has proved compelling; it is certainly the most commonly mentioned explanation for this particular pattern of trading.

When researchers make a link between prospect theory and the disposition effect, they typically do so using an informal argument. To be sure that prospect theory really does predict a disposition effect, and therefore that it is a valid hypothesis for further study, some formal modelling is needed. To date, there has been very little such work.

In this paper, we take up this task, and present a simple but rigorous model of trading behavior for an investor with prospect theory preferences. Specifically, we consider an investor who, at the beginning of the year, buys shares of a stock. Over the course of the year, she trades the stock, and, at the end of the year, receives prospect theory utility based on her 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 investor's optimal trading strategy. This allows us to simulate artificial data on how prospect theory investors would trade over time, and 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 two main findings. First, for *some* values of μ and T, prospect theory does indeed predict a disposition effect. As such, our model offers a formalization of the intuitive arguments that have been used in the past to link prospect theory and the disposition effect. Our second result is more surprising. We find that for other, equally reasonable values of μ and T, the standard intuition breaks down, and prospect theory predicts the *opposite* of the disposition effect: that investors will be more inclined to sell stocks with prior *losses* than stocks with prior gains.

We demonstrate this last result in detail in Section 3. The basic idea, however, can be illustrated with a two-period example. In the two-period case, and for the preference parameters estimated by Tversky and Kahneman (1992), our implementation of prospect theory always predicts the opposite of the disposition effect, whatever the expected stock return μ.

To see the intuition, suppose that a prospect theory investor buys a share of a stock for $50. Since the investor is loss averse, the fact that she bought the stock at all means that it must have a high expected return: that, for example, it will go up by $10 or fall by $5 each period, with equal chance.

Suppose that, over the first period, the stock rises $10 to $60. Our analysis shows that the investor will now take a position in the stock such that, even if the stock does poorly in the next period, she will still just about break even. In other words, she will now hold approximately *two* shares of the stock, because even if the stock falls $5 to $55, she will still

break even overall:

$$\$10 + 2(\$55 - \$60) = \$0. \tag{1}$$

The intuition behind this strategy is that, based on the estimates of Tversky and Kahneman (1992), the prospect theory value function is only *mildly* concave over gains. The investor is therefore almost risk-neutral in this region, and is willing to risk the loss of her initial gain.

Now suppose that, over the first period, the stock falls \$5 to \$45. Our analysis shows that the investor will now take a position in the stock such that, if the stock does well in the next period, she will again just about break even. In other words, she will now hold approximately 0.5 shares of the stock, because if the stock rises \$10 to \$55, she will indeed break even:

$$-\$5 + 0.5(\$55 - \$45) = 0. \tag{2}$$

The intuition behind this strategy is that, since the prospect theory value function is convex over losses, the investor is willing to gamble just enough to give herself a chance of making back her initial loss.

This simple example illustrates the mechanism driving the result that we demonstrate more rigorously in Section 3: that, in its simplest implementation, prospect theory often predicts the opposite of the disposition effect. After an initial gain, the investor increases her allocation to two shares; after a loss, she reduces her allocation to 0.5 shares. She therefore sells after a *loss* rather than after a gain. The pitfall in the traditional argument linking prospect theory to the disposition effect is that it does not take into account the investor's *initial* buying decision. As soon as we do, we realize that the expected return on the stock must be high, and hence that the size of the initial gain (\$10) must exceed the size of the initial loss (\$5). It therefore takes a larger share allocation to break even after a gain (equation (1)) than it does to break even after a loss (equation (2)).

Our results are relevant not only for the trading of individual investors in the stock market, but also for the disposition-type effects that have been documented in settings as varied as the housing market, futures trading, and executive stock options; as well as for the recent evidence linking the disposition effect to momentum and post-earnings announcement drift in stock returns. All of these findings have been linked to prospect theory. In light of our results, however, we may need to go back and check that the links can be formally justified.

The conclusion we draw from our analysis is that, for one, simple implementation of prospect theory, the connection between these preferences and the disposition effect is not as obvious as previously thought. At the same time, we are keen to emphasize that this implementation *can*, sometimes, generate a disposition effect. As such, it may turn out to be the right way of thinking about this particular pattern of trading. To determine this, future research could test some of the new predictions, discussed in Section 3, that come out

of our model. If these predictions are not borne out in the data, we can turn our attention to other models of the disposition effect, including, for example, other implementations of prospect theory.

In Section 2, we review the disposition effect, prospect theory, and the argument that has been used to link the two. In Section 3, we analyze the trading behavior of an investor with prospect theory preferences, and use our model to see if prospect theory generates a disposition effect. In Section 4, we discuss the robustness of our results and their implications for a number of recent empirical findings. Section 5 concludes.

2   The Disposition Effect: Evidence and Interpretation

Odean (1998) analyzes the trading activity, from 1987 to 1993, of 10,000 households with accounts at a large discount brokerage firm. He finds that, when an investor in his sample sells shares, she prefers to sell shares of a stock that has *risen* in value since purchase than of one that has fallen in value. Specifically, for any day on which an investor in the sample sells shares of a stock, 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, the ratio PGR is computed:

$$\text{PGR} = \frac{\text{no. of realized gains}}{\text{no. of realized gains} + \text{no. of paper gains}}. \tag{3}$$

In words, PGR ("Proportion of Gains Realized") 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,

$$\text{PLR} = \frac{\text{no. of realized losses}}{\text{no. of realized losses} + \text{no. of paper losses}}, \tag{4}$$

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 though this effect is, its cause remains unclear. Many of the most obvious potential explanations fail to capture important features of the data. Perhaps the most obvious hypothesis of all is the information hypothesis: investors sell stocks with paper gains because they have private information that these stocks will subsequently do poorly, and they 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 inclination to sell stocks with paper *losses* because the losses thus realized can be used to offset taxable gains in other assets.[1]

Odean (1998) also casts doubt on the hypothesis that the disposition effect is nothing more than portfolio rebalancing. 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.

Finally, Odean (1998) finds little evidence for the idea that investors refrain from selling past losers because these stocks, by virtue of having fallen in value, often trade at low prices, where transaction costs can be higher. Specifically, he shows that PGR continues to exceed PLR even when the sample is restricted to stocks whose prices exceed $10 per share, a range for which differences in transaction costs between prior winners and losers are small.

Given the difficulties faced by these standard hypotheses, two alternative explanations have been proposed. One is that, for some reason, individual investors have an irrational belief in mean-reversion: even though the prior winners they sell subsequently outperform the prior losers they hold on to, investors *think* that the prior winners will underperform and that the prior losers will outperform.

The other non-standard explanation is based on Kahneman and Tversky's (1979) prospect theory, a leading theory of decision-making under risk. Before discussing this explanation in more detail, we briefly review the main features of prospect theory.

2.1 Prospect theory

Consider the gamble

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

to be read as "get 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 utility framework, an agent with utility function $U(\cdot)$ evaluates this risk by computing

$$pU(W + x) + qU(W + y), \tag{5}$$

---

[1] 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 deadline for realizing losses approaches.

where W is her current wealth. In the framework of prospect theory, the agent assigns the gamble the value

$$\pi(p)v(x) + \pi(q)v(y), \tag{6}$$

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 (5) and (6). First, the carriers of value in prospect theory are gains and losses, not final wealth levels: the argument of $v(\cdot)$ in (6) is x, not $W + x$. This is motivated in part by experimental evidence, but is also consistent with the way in which our perceptual apparatus is more attuned to a *change* in the level of an attribute – brightness, loudness, or temperature, say – than to the level itself.

Second, the value function $v(\cdot)$ is concave over gains, but convex over losses. Kahneman and Tversky (1979) infer this from subjects' preference for a certain gain of \$500 over[2]

$$(\$1,000, \frac{1}{2}),$$

and from their preference for

$$(-\$1,000, \frac{1}{2})$$

over a certain loss of \$500. In short, people are risk averse over moderate-probability gains, but risk-seeking over moderate-probability losses.

Third, the value function 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 element of prospect theory is known as loss aversion. Kahneman and Tversky (1979) infer the kink from the widespread aversion to bets of the form

$$(\$110, \frac{1}{2}; -\$100, \frac{1}{2}). \tag{7}$$

Such aversion is hard to explain with differentiable utility functions, whether expected utility or non-expected utility, because the very high local risk aversion required to do so typically predicts implausibly high aversion to large-scale gambles (Epstein and Zin, 1990; Rabin, 2000; Barberis, Huang, and Thaler, 2006).

Finally, under prospect theory, the agent does not use objective probabilities when evaluating a gamble, but rather, transformed probabilities obtained from objective probabilities via the probability weighting function $\pi(\cdot)$. The primary effect of this function is to overweight low probabilities, a feature that parsimoniously captures the simultaneous demand many individuals have for both lotteries *and* insurance.

---

[2]We abbreviate $(x, p; 0, q)$ to (x, p).

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

$$v(x) = \begin{cases} x^\alpha \\ -\lambda(-x)^\alpha \end{cases} \text{for} \quad \begin{matrix} x \geq 0 \\ x < 0 \end{matrix}, \ 0 < \alpha < 1, \lambda > 1. \tag{8}$$

For $0 < \alpha < 1$ 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 $\alpha = 0.88$ and $\lambda = 2.25$. Figure 1 plots the function in (8) for these parameter values. An α of 0.88 means that the value function is only mildly concave over gains and only mildly convex over losses, while a λ of 2.25 makes the agent substantially more sensitive to losses than to gains. This will turn out to be important in what follows.[3]

2.2 Prospect theory and the disposition effect

Over the past 20 years, many papers have drawn a connection between prospect theory and the disposition effect. Shefrin and Statman (1985) was the first paper to do so; Weber and Camerer (1998), Odean (1998), and Grinblatt and Han (2005) are just a few of the articles that followed. Most of the time, the link between prospect theory and the disposition effect is described in informal terms. The argument varies slightly from paper to paper, but the essence is always the same: A stock that has risen in value since purchase brings the investor into the concave, risk averse, "gain" region of the value function in Figure 1. A stock that has fallen in value since purchase brings the investor into the convex, risk-seeking, "loss" region of the value function. As a result, the agent is more willing to take risk on a stock trading at a paper loss, than on one trading at a paper gain. She will therefore be more inclined to sell the latter.

The prospect theory view of the disposition effect has proved compelling: it is the most commonly mentioned explanation for this pattern of trading. To be sure that prospect theory really does predict a disposition effect, and therefore that it is a valid hypothesis for further study, some formal modelling is needed. In Section 3, we take up this task, and present a simple but rigorous model of trading behavior for an investor with prospect theory preferences.

---

[3]Strictly speaking, the value function in equation (8) 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, it makes the agent more sensitive to losses than to gains, and explains the rejection of bets like that in (7).

3 The Trading Behavior of Agents with Prospect Theory Preferences

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 each period, and a risky asset, which we think of as an individual stock. The price 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:

$$R_{t,t+1} = \begin{cases} R_u > R_f & \text{with probability } \pi \\ R_d < R_f & \text{with probability } 1 - \pi \end{cases} , \text{ i.i.d. across periods,} \tag{9}$$

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

$$\pi R_u + (1 - \pi)R_d > R_f, \tag{10}$$

so that the expected stock return exceeds the risk-free rate.

We study the trading behavior of an investor with prospect theory preferences, who, in particular, uses the value function $v(\cdot)$ in equation (8). The argument of $v(\cdot)$ is the investor's "gain" or "loss". Prospect theory does not specify exactly what the gain or loss should be. In our context, the simplest approach is to define the gain or loss as the profit from trading the stock over the interval from 0 to T; in symbols, as

$$\Delta W_T \equiv W_T - W_0, \tag{11}$$

where W_t is the investor's wealth at time t. In this paper, we work with a slightly adjusted version of (11), namely

$$\Delta W_T = W_T - W_0 R_f^T, \tag{12}$$

so that the investor defines her gain or loss as her trading profit over the interval from 0 to T, relative to the profit she 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 her trading a success if it earns her 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.[4]

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 skewed securities, which deliver a very

---

[4]Researchers have not, as yet, been able to pinpoint exactly which reference levels people use in practice. We therefore choose a reference level – initial wealth scaled up by the risk-free rate – that is simple and tractable, and that can serve as a benchmark for future analysis.

good or very bad outcome with low probability. Since most stocks are not highly skewed, we focus mainly on values of π close to 0.5, where probability weighting has a negligible effect.[5]

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

$$\max_{x_0, x_1, \ldots, x_{T-1}} E[v(\Delta W_T)] = E[v(W_T - W_0 R_f^T)], \tag{13}$$

where $v(\cdot)$ is defined in equation (8), subject to the budget constraint

$$\begin{aligned} W_t &= (W_{t-1} - x_{t-1} P_{t-1}) R_f + x_{t-1} P_{t-1} R_{t-1,t} \\ &= W_{t-1} R_f + x_{t-1} P_{t-1} (R_{t-1,t} - R_f), \qquad t = 1, \ldots, T, \end{aligned} \tag{14}$$

and a non-negativity of wealth constraint

$$W_T \geq 0. \tag{15}$$

By taking the investor's gain or loss to be the profit earned from trading an *individual* stock, we are following the informal arguments that have been used in the literature to link prospect theory with the disposition effect. This assumption means that the investor engages in what is sometimes called "narrow framing" or "mental accounting," in other words, that she gets utility *directly* from the outcome of her investment in a single stock, even if this is just one of many stocks in her overall portfolio. While this is not a standard assumption in finance models, we do not take a stand on its plausibility here. Our goal is simply to investigate whether, as has been suggested using informal arguments, prospect theory in combination with narrow framing predicts a disposition effect.[6]

When we parameterize our model in Section 3.1, we take the interval from 0 to T to be one year. An informal summary of our framework is therefore that, at the start of the year, time 0, the investor buys some shares of a stock, and then trades the stock over the course of the year. At the end of the year, time T, she receives prospect theory utility defined over her trading profit.

An alternative approach is to posit that the investor receives utility whenever she sells shares of the stock. If she sells shares at time $t < T$, say, she receives a jolt of utility at time t based on the size of the realized gain or loss: positive utility if a gain is realized, and negative otherwise. This approach is more radical than the one we study here, in that it

---

[5]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.

[6]Since we do think of the risky asset as an individual stock, initial wealth W_0 is best interpreted as the maximum amount the investor is willing to lose from trading the stock.

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.

We think that, methodologically, it makes sense to first study the model with the fewer deviations from the traditional framework, in other words, the model in (13)-(15). If it turns out that this model does predict a disposition effect, we will learn that there is no need to appeal to utility from realized gains in order to understand this pattern of trading. If the model in (13)-(15) turns out not to predict a disposition effect, we will then have license to move on and investigate the more radical model.[7]

We now present two approaches to solving the problem in (13)-(15). Naturally, the two approaches lead to the same optimal allocations x_t. The reason we present both is that each has important advantages. First, in Section 3.1, we present an approach due to Cox and Huang (1989) whereby, when markets are complete, the dynamic problem in (13)-(15) can be rewritten as a simpler static problem. This approach is very powerful: for any number of trading periods, it gives a fully analytical solution for the investor's trading strategy. Its drawback is that it does not offer much intuition for why the results turn out the way they do. In Section 3.2, we therefore present an alternative approach, based on standard dynamic programming techniques. This approach becomes computationally intractable as the number of time periods grows, but it provides a great deal of intuition.

3.1   A complete markets approach

Cox and Huang (1989) demonstrate that, when markets are complete, an investor's dynamic optimization problem can be rewritten as a *static* problem in which the investor directly chooses her wealth in the different possible states at the final date. An optimal trading strategy is one which 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, $j = 1, 2, \cdots, t + 1$, 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 therefore $P_0 R_u^{t-j+1} R_d^{j-1}$.

We denote the investor's optimal share allocation in node j at time t as $x_{t,j}$; the optimal

---

[7]Since the earlier papers linking prospect theory and the disposition effect typically use informal arguments to make the link, they often fail to specify exactly how the "gains" and "losses" of prospect theory should be defined. As a result, it is sometimes hard to tell which of the two models they have in mind. A close reading of these earlier papers suggests that about half implicitly have the model in (13)-(15) in mind, while the remainder are implicitly appealing to utility from realized gains and losses.

wealth in that node as $W_{t,j}$; and the ex-ante probability of reaching that node as $\pi_{t,j}$, so that

$$\sum_{j=1}^{t+1} \pi_{t,j} = 1. \tag{16}$$

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

$$q_{t,j} = p_{t,j}/\pi_{t,j}. \tag{17}$$

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

$$\sum_{j=1}^{t+1} \pi_{t,j} q_{t,j} = \frac{1}{(R_f)^t}. \tag{18}$$

With this notation in hand, we apply Cox and Huang's (1989) insight and rewrite the dynamic optimization problem in (13)-(15) as

$$\max_{\{W_{T,j}\}_{j=1,\ldots,T+1}} \sum_{j=1}^{T+1} \pi_{T,j} v(W_{T,j} - W_0 R_f^T), \tag{19}$$

subject to the budget constraint

$$\sum_{j=1}^{T+1} \pi_{T,j} q_{T,j} W_{T,j} = W_0 \tag{20}$$

and a non-negativity of wealth constraint

$$W_{T,j} \geq 0, \qquad j = 1, \ldots, T+1. \tag{21}$$

This static problem can be solved using Lagrange multiplier techniques. We summarize the solution in Proposition 1. 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. In the proof of the proposition, we show that, under this assumption, the ex-ante probability of reaching node j at time t, $\pi_{t,j}$, is given by

$$\pi_{t,j} = \frac{t! 2^{-t}}{(t-j+1)!(j-1)!}, \qquad 0 \leq t \leq T, \ 1 \leq j \leq t+1; \tag{22}$$

and that $q_{t,j}$, the state price density in node j at time t, is given by

$$q_{t,j} = q_u^{t-j+1} q_d^{j-1}, \qquad 0 \leq t \leq T, \ 1 \leq j \leq t+1, \tag{23}$$

where

$$q_u = \frac{2(1-R_d)}{R_u - R_d}, \qquad q_d = \frac{2(R_u - 1)}{R_u - R_d}, \tag{24}$$

12

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*

$$V^* = \max_{k \in \{1,\dots,T\}} \left[\left(\sum_{l=1}^{k} q_{T,l}^{-\frac{\alpha}{1-\alpha}} \pi_{T,l} \right)^{1-\alpha} \left(\sum_{l=k+1}^{T+1} q_{T,l} \pi_{T,l} \right)^{\alpha} - \lambda \sum_{l=k+1}^{T+1} \pi_{T,l} \right], \qquad (25)$$

and let k^* be the $k \in \{1,\dots,T\}$ at which the maximum in (25) is attained.

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

$$W_{T,j} = \begin{cases} W_0 R_f^T \left[1 + q_{T,j}^{-\frac{1}{1-\alpha}} \frac{\sum_{l=k^*+1}^{T+1} q_{T,l} \pi_{T,l}}{\sum_{l=1}^{k^*} q_{T,l}^{-\frac{\alpha}{1-\alpha}} \pi_{T,l}} \right] & \text{if } j \leq k^* \\ 0 & \text{if } j > k^* \end{cases} \qquad (26)$$

if $V^* > 0$; and by
$$W_{T,j} = W_0 R_f^T, \qquad j = 1, \dots, T+1, \qquad (27)$$

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

$$x_{t,j} = \frac{W_{t+1,j} - W_{t+1,j+1}}{P_0(R_u^{t-j+2} R_d^{j-1} - R_u^{t-j+1} R_d^j)}, \qquad 0 \leq t \leq T, \ 1 \leq j \leq t+1, \qquad (28)$$

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

$$W_{t,j} = \frac{\frac{1}{2} W_{t+1,j} q_{t+1,j} + \frac{1}{2} W_{t+1,j+1} q_{t+1,j+1}}{q_{t,j}}, \qquad 0 \leq t \leq T, \ 1 \leq j \leq t+1. \qquad (29)$$

Proof of Proposition 1. 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 (26) and (27). We find that the investor's optimal policy is either to choose an allocation equal to the reference wealth level $W_0 R_f^T$ in all date T nodes, as in (27); or, as in (26), to use a "threshold" strategy, in which, for some $k^* : 1 \leq k^* \leq T$, she allocates a wealth level greater than the reference level $W_0 R_f^T$ 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 (25) maximizes the investor's utility across the T possible values of k^*. If the best threshold strategy offers negative utility, $V^* < 0$, which occurs when the expected risky asset return is low, the investor does not use a threshold strategy, and instead chooses a wealth level of $W_0 R_f^T$ in all final date nodes; otherwise, she adopts the best threshold strategy.

13

We now illustrate the proposition with a specific 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: a year, say. 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

$$\left(\frac{R_u + R_d}{2}\right)^T = \mu, \qquad \left(\frac{R_u^2 + R_d^2}{2}\right)^T = \mu^2 + \sigma^2, \tag{30}$$

which imply

$$R_u = \mu^{\frac{1}{T}} + \sqrt{(\mu^2 + \sigma^2)^{\frac{1}{T}} - (\mu^2)^{\frac{1}{T}}} \tag{31}$$

$$R_d = \mu^{\frac{1}{T}} - \sqrt{(\mu^2 + \sigma^2)^{\frac{1}{T}} - (\mu^2)^{\frac{1}{T}}}. \tag{32}$$

In our example, we set $(\mu, \sigma) = (1.1, 0.3)$, which, from (31)-(32), corresponds to $(R_u, R_d) = (1.16, 0.89)$.

For these parameter values, the top-left panel in Table 1 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 (23) and (24). 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 the proposition: 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 tends to take more risk after a *gain* in the stock than after a loss; this behavior is the opposite of the disposition effect.

We now investigate more carefully whether prospect theory predicts a disposition effect. In brief, we use Proposition 1 to simulate an artificial dataset of how prospect theory investors would trade over time. We mimic, as much as possible, the structure of Odean's (1998) actual dataset. We then apply Odean's (1998) methodology to see if, in our simulated data, investors exhibit a disposition effect.

Odean's (1998) data cover 10,000 accounts. We therefore generate trading data for 10,000 investors with prospect theory preferences, each of whom holds N_S stocks. For each investor, we use the binomial distribution in (9) to simulate a T-period stock price path for each of her N_S stocks. We assume that all stocks have the same annual expected return μ and

14

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 dataset of how the 10,000 prospect theory investors trade each of their N_S stocks over 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 1, we know that the investor will trade the stock so as to allocate 1.7, 1.8, 3.5, and 0.5 shares at each trading date.

To see if there is a disposition effect in our artificial data, we follow the method of Odean (1998), described in Section 2. 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 her stocks at date $t \in \{1, \ldots, T - 1\}$, we count a "realized gain" if the stock price exceeds the average price at which shares were purchased, and a "realized loss" otherwise. For every stock in the investor's portfolio at date t that is *not* sold, we count a "paper gain" if the stock price exceeds the average price at which 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 compute PGR and PLR, first defined in equations (3)-(4):

$$\text{PGR} = \frac{\text{no. of realized gains}}{\text{no. of realized gains} + \text{no. of paper gains}} \tag{33}$$

$$\text{PLR} = \frac{\text{no. of realized losses}}{\text{no. of realized losses} + \text{no. of paper losses}}. \tag{34}$$

We say, as does Odean (1998), 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 households in his sample, but Barber and Odean (2000), who use very similar data, report a mean value slightly above 4. We therefore set $N_S = 4$. Our results are relatively insensitive to the value of N_S.

Table 2 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, we simulate an artificial dataset and use Odean's (1998) methodology to compute PGR and PLR. The boldface type identifies cases where PGR is less than PLR; in words, cases where the disposition effect fails. 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 hyphens. 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, the kink in the utility function at time T

is smoothed out, lowering investors' initial risk aversion, and increasing their willingness to buy the risky asset.

The table illustrates the main results of the paper. First, we see that, in many cases, our simple implementation of prospect theory does predict a disposition effect: PGR often exceeds PLR. As such, our model offers a formalization of the informal arguments that have been used to link prospect theory and the disposition effect.

Our second result is more surprising. We find that, in many cases, PGR is *lower* than PLR. Put differently, prospect theory often predicts the *opposite* of the disposition effect, namely that investors prefer to sell a stock trading at a paper *loss* than one trading at a paper gain. Moreover, the table 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 disposition effect always fails, while for $T = 12$, it fails in about half the cases we report. In the next section, we try to understand what is driving these results.[8]

3.2   A dynamic programming approach

The Cox-Huang (1989) approach of Section 3.1 is very powerful: for any number of trading periods, it provides an analytical solution for optimal wealth allocations and share holdings. Its one drawback is that it offers little intuition for *why* the results turn out the way they do: Why, in many cases, does the disposition effect fail? Why does it tend to fail when the expected risky asset return is high or when the number of trading periods is low? To address these questions, we turn to an alternative method for solving problem (13)-(15), namely dynamic programming. This approach complements the Cox-Huang (1989) method. While it does not lead to an analytical solution and is not computationally tractable when the number of trading periods is large, it offers much more intuition.

Since our goal in this section is to provide intuition, we keep things simple by setting $T = 2$, 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 especially instructive because, as shown in Table 2, the disposition effect always fails in this case, at least for the preference parameters estimated by Tversky and Kahneman (1992). We also set the gross risk-free rate to $R_f = 1$.

---

[8]It is hard to know which values of μ and T are most reasonable. We therefore present results for a range of values. For some readers, the most reasonable values of μ and T may be those corresponding 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 which has not been noted in the literature to date: that, for *some* parameters, prospect theory can predict the opposite of the disposition effect.

With these simplifications, the decision problem (13)-(15) becomes

$$\max_{x_0,x_1} E[v(\Delta W_2)] \quad = \quad E[v(W_2 - W_0)] = E\left[v\left(x_0 P_0(R_{0,1} - 1) + x_1 P_1(R_{1,2} - 1)\right)\right] \quad (35)$$
$$\text{subject to } W_2 \quad \geq \quad 0,$$

where the expression for $W_2 - W_0$ comes from combining the budget constraints

$$W_1 = W_0 + x_0 P_0(R_{0,1} - 1), \qquad W_2 = W_1 + x_1 P_1(R_{1,2} - 1). \qquad (36)$$

In the dynamic programming approach, we first solve the time 1 decision problem, and then work backwards to find the optimal share holdings at time 0. At time 1, the quantity $x_0 P_0(R_{0,1} - 1)$ in (35) is known, and is therefore a state variable for the time 1 value function. Since

$$x_0 P_0(R_{0,1} - 1) = W_1 - W_0 \equiv \Delta W_1, \qquad (37)$$

we see that $x_0 P_0(R_{0,1} - 1)$ is the investor's gain or loss from trading between time 0 and time 1, and refer to it as the time 1 gain/loss, ΔW_1. At time $t = 1$, then, the investor solves

$$J(\Delta W_1) \quad \equiv \quad \max_{x_1} E[v(\Delta W_1 + x_1 P_1(R_{1,2} - 1))] \qquad (38)$$
$$\text{subject to } W_2 \quad \geq \quad 0,$$

and at time $t = 0$, solves

$$\max_{x_0} E[J(x_0 P_0(R_{0,1} - 1))]$$
$$\text{subject to } W_1 \geq 0. \qquad (39)$$

In the two-period setting of this section, we say that there is a disposition effect if and only if

$$x_{1,1} < x_{0,1} \leq x_{1,2}, \qquad (40)$$

where, as in Section 3.1, $x_{t,j}$ is share holdings in the j'th node from the top at date t in the binomial tree that describes the evolution of the stock price. In words, condition (40) says that there is a disposition effect if the investor sells shares after a time 1 gain ($x_{1,1} < x_{0,1}$) and buys shares or maintains the same position after a time 1 loss ($x_{0,1} \leq x_{1,2}$), so that the relative propensity to sell shares is greater after a gain.

Condition (40) is consistent with the definition of the disposition effect in Section 3.1 and in Odean (1998), namely that PGR > PLR. For example, if we have data on the trading activity of a large number of investors, each of whom trades one stock for two periods, say, and we compute PGR and PLR using Odean's (1998) methodology, then it is straightforward to see that

$$\begin{array}{llll} \text{PGR} & > & \text{PLR} & \text{if } x_{1,1} < x_{0,1} \leq x_{1,2} \\ \text{PGR} & < & \text{PLR} & \text{if } x_{1,2} < x_{0,1} \leq x_{1,1} \\ \text{PGR} & = & \text{PLR} & \text{otherwise.} \end{array}$$

Condition (40) is therefore a natural definition of the disposition effect in a two-period setting.[9]

In Table 2, we saw that, in a two-period setting $(T = 2)$, and for the preference parameterization estimated by Tversky and Kahneman (1992), our implementation of prospect theory always predicts the opposite of the disposition effect. We now quickly check that this result remains true when we use the definition of the disposition effect in condition (40).

Specifically, we solve problems (38) and (39) for various parameter values and check whether condition (40) is satisfied. 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 each μ, Table 3 reports optimal share holdings at time 0, $x_{0,1}$, time 1 share holdings after a gain, $x_{1,1}$, and time 1 share holdings after a loss, $x_{1,2}$. The table confirms the results in the "$T = 2$" column of Table 2. When the expected risky asset return is below 1.1, loss aversion prevents the investor from buying any stock at all; for higher values of the expected return – in other words, for values where we can check for a disposition effect – the disposition effect always fails: the investor always sells after a loss and buys after a gain, so that condition (40) is violated.

To understand *why* the disposition effect fails in the two-period case, we first study the time 1 share holdings that solve problem (38), $x_1(\Delta W_1)$. Indeed, the reason why the dynamic programming approach offers more intuition than the Cox-Huang (1989) approach is that the former tells us the optimal time 1 share holdings for *any* time 1 gain/loss ΔW_1, while the latter only tells us the time 1 share holdings that correspond to one specific time 1 gain/loss, namely the optimal one.

Proposition 2. *The investor's time 1 share holdings $x_1(\Delta W_1)$ depend on the parameter*

$$g = (R_u - 1)/(1 - R_d). \tag{41}$$

- *When $0 < g < \lambda^{\frac{1}{\alpha}}$, optimal share holdings are given by*

$$x_1(\Delta W_1) = \begin{cases} \min\left[x_L(\Delta W_1), \frac{W_0 + \Delta W_1}{P_1(1 - R_d)}\right] & \Delta W_1 < 0 \\ 0 & \text{for } \Delta W_1 = 0 \\ x_G(\Delta W_1) & \Delta W_1 > 0 \end{cases}, \tag{42}$$

---

[9]Initially, it may seem that $x_{1,1} < x_{1,2}$ is also a reasonable definition of the disposition effect in a two-period setting: the investor holds less of the stock after a gain than after a loss. However, this condition is not consistent with Odean's (1998) definition, that PGR > PLR. If $x_{1,1} < x_{1,2} < x_{0,1}$, for example, the investor realizes both gains and losses, so that PGR = PLR. And if $x_{0,1} \leq x_{1,1} < x_{1,2}$, the investor realizes neither gains nor losses, so that, once again, PGR = PLR. Only condition (40) is consistent with PGR > PLR.

where

$$x_L(\Delta W_1) \;=\; \frac{-\Delta W_1}{P_1} \left(\frac{R_u - R_d}{\left(\frac{R_u-1}{\lambda(1-R_d)} \right)^{\frac{1}{1-\alpha}} + 1} - (1 - R_d) \right)^{-1} \tag{43}$$

$$x_G(\Delta W_1) \;=\; \frac{\Delta W_1}{P_1} \left((1 - R_d) + \frac{R_u - R_d}{\left(\frac{R_u-1}{1-R_d} \right)^{\frac{1}{1-\alpha}} - 1} \right)^{-1}. \tag{44}$$

- *When $g > \lambda^{\frac{1}{\alpha}}$, optimal share holdings are given by*

$$x_1(\Delta W_1) = \begin{cases} \frac{W_0 + \Delta W_1}{P_1(1-R_d)} & \\ x_G(\Delta W_1) & \end{cases} \quad \text{for} \quad \begin{array}{c} \Delta W_1 \leq \Delta W^* \\ \Delta W_1 > \Delta W^* \end{array}, \tag{45}$$

where ΔW^* is the unique positive solution to

$$\left[(g+1)\frac{\Delta W^*}{W_0} + g \right]^{\alpha} - \left(\frac{1+g}{1+g^{\frac{\alpha}{1-\alpha}}} \right)^{\alpha} \left(g^{\frac{\alpha^2}{1-\alpha}} + g^{-\alpha} \right) \left(\frac{\Delta W^*}{W_0} \right)^{\alpha} = \lambda. \tag{46}$$

Proof of Proposition 2. See the Appendix.

Equations (42) and (45) make reference to the quantity $(W_0 + \Delta W_1)/(P_1(1 - R_d))$. This is the largest share allocation allowed by the constraint that time 2 wealth be non-negative. To see this, note that since

$$W_2 = W_1 + x_1 P_1(R_{1,2} - 1), \tag{47}$$

we need $x_1 \leq W_1/(P_1(1 - R_d))$ to ensure that W_2 remains non-negative even if the return from time 1 to time 2 is R_d. Since, by definition, $W_1 = W_0 + \Delta W_1$, this implies

$$x_1 \leq \frac{W_0 + \Delta W_1}{P_1(1 - R_d)}. \tag{48}$$

The proposition distinguishes between two cases, depending on the level of g. To see why prospect theory fails to predict a disposition effect, we need to consider each case in turn. The case of $g > \lambda^{\frac{1}{\alpha}}$ is straightforward. In this case, the expected return on the risky asset is high. The investor finds the risky asset so attractive that, as equation (45) shows, for much of the range of ΔW_1, her allocation is limited only by the wealth constraint. After a time 1 loss, the investor has relatively little wealth, and is forced by the wealth constraint to take a small position in the risky asset. After a gain, she is wealthier, and, even if wealth constrained, can take a larger position. She therefore takes more risk after a gain than after a loss, contrary to the disposition effect. This is why, for $\mu \geq 1.18$ in Table 3 – the range for which $g > \lambda^{\frac{1}{\alpha}}$ – the disposition effect fails.

19

The more interesting case is that of $g < \lambda^{\frac{1}{\alpha}}$. Here, the expected risky asset return is relatively low. We assume, however, that the expected return, while low, is still somewhat higher than the risk-free rate; as we will see in the next section, this is the only case that is relevant to understanding whether the disposition effect holds. For an expected return that is low but not too low, then, the investor adopts the following strategy. After a gain at time 1 ($\Delta W_1 > 0$), she takes a position such that, even if the return from time 1 to time 2 is poor, she still ends up with a gain – a small gain – at time 2.

This result can be seen both mathematically and intuitively. To see it mathematically, look at equation (44), which, from equation (42), gives the optimal share allocation after a time 1 gain. So long as the expected risky asset return is not too low, the quantity $(\frac{R_u-1}{1-R_d})^{\frac{1}{1-\alpha}}$ is relatively large, and so $x_1(\cdot)$ is slightly less than $\frac{\Delta W_1}{P_1(1-R_d)}$. Since the investor's gain/loss at time 2 is

$$\Delta W_2 = \Delta W_1 + x_1 P_1(R_{1,2} - 1),$$

we see that, at time 1, the investor takes a position in the risky asset such that, if the return from time 1 to time 2 is poor – in other words, if $R_{1,2} = R_d$ – she ends up with a time 2 gain slightly greater than

$$\Delta W_1 + \frac{\Delta W_1}{P_1(1 - R_d)} P_1(R_d - 1) = 0; \tag{49}$$

in words, a time 2 gain slightly greater than zero.

What is the intuition for this result? The concavity of the value function $v(\cdot)$ in the region of gains is mild – Tversky and Kahneman's (1992) estimate of $\alpha = 0.88$ implies only mild concavity. Moreover, we are focussing on the case where the expected risky asset return, while low, is still somewhat higher than the risk-free rate. Taken together, these two things mean that the investor is willing to gamble at least as far as the edge of the concave region; in other words, to take a time 1 position such that, after a poor return from time 1 to time 2, her time 2 gain is slightly greater than zero. However, she is not willing to take a larger gamble than this: if she does, she risks ending up with a loss at time 2, which, given that she is loss averse, would be very painful.

If the investor has a *loss* at time 1 ($\Delta W_1 < 0$), she will, at most, take a position such that, if the return from time 1 to time 2 is good, she ends up with a time 2 gain slightly above zero. Once again, this can be understood both mathematically and intuitively.

For a mathematical perspective, look at equation (43), which, from equation (42), is an upper bound on the investor's optimal share allocation after a loss. Since $g < \lambda^{\frac{1}{\alpha}}$, the quantity $(\frac{R_u-1}{\lambda(1-R_d)})^{\frac{1}{1-\alpha}}$ is small, and so $x_1(\cdot)$ is slightly greater than $\frac{-\Delta W_1}{P_1(R_u-1)}$. This means that, at time 1, the investor takes a position in the risky asset such that, if the return from time 1 to time 2 is good, she ends up with a time 2 gain slightly greater than

$$\Delta W_1 + \frac{-\Delta W_1}{P_1(R_u - 1)} P_1(R_u - 1) = 0; \tag{50}$$

20

in words, a time 2 gain slightly greater than zero.

The intuition for this is that, since the value function $v(\cdot)$ is convex in the region of losses, the investor is happy to gamble at least as far as the edge of the convex region; in other words, to take a time 1 position such that, after a good return from time 1 to time 2, she ends up with a time 2 gain close to zero. However, she is not willing to take a position much larger than this because, to the right of the kink in the value function, the potential marginal gain is much lower relative to the potential marginal loss.

3.3 An example

We now illustrate the discussion in Section 3.2 graphically; once we have done so, it will be easier to see why prospect theory fails to predict a disposition effect in the case of $g < \lambda^{\frac{1}{\alpha}}$. In the example we now consider, we set $(P_0, W_0) = (40, 40)$, $(\mu, \sigma) = (1.1, 0.3)$, and $(\alpha, \lambda) = (0.88, 2.25)$, which correspond to the third row of Table 3. When $T = 2$, this choice of μ and σ implies $(R_u, R_d) = (1.25, 0.85)$. For these parameters, g is indeed less than $\lambda^{\frac{1}{\alpha}}$.

Given these parameters, we can compute the time 1 allocation function $x_1(\Delta W_1)$ in equation (42), the time 1 value function $J(\cdot)$ in (38), and hence the time 0 allocation x_0 that solves problem (39). We find that $x_0 = 4.0$. Figure 2 plots the investor's time 1 and time 2 gain/losses on her time 2 prospect theory utility function. Point A marks her potential time 1 gain, namely

$$\Delta W_1 = x_0 P_0 (R_u - 1) = (4.0)(40)(0.25) = 39.9.$$

From equation (44), the corresponding optimal time 1 allocation is $x_1(39.9) = 5.05$. Points B and B' mark the time 2 gain/losses that this time 1 allocation could lead to, namely

$$\Delta W_1 + x_1 P_1 (R_u - 1) = 39.9 + (5.05)(40)(1.25)(0.25) = 102.71, \text{ or}$$
$$\Delta W_1 + x_1 P_1 (R_d - 1) = 39.9 + (5.05)(40)(1.25)(-0.15) = 1.63.$$

Similarly, Point C marks the investor's potential time 1 loss, namely

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

From equation (42), the corresponding optimal time 1 allocation is $x_1(-24.3) = 3.06$. Points D and D' mark the time 2 gain/losses that this time 1 allocation could lead to, namely[10]

$$\Delta W_1 + x_1 P_1 (R_u - 1) = -24.3 + (3.06)(40)(0.85)(0.25) = 1.63, \text{ or}$$
$$\Delta W_1 + x_1 P_1 (R_u - 1) = -24.3 + (3.06)(40)(0.85)(-0.15) = -40.$$

---

[10]B and D' are the same point. The proof of Proposition 1 shows that our model satisfies a "path independence" property, whereby the optimal gain/loss at any date T node is independent of the path the stock price takes through the binomial tree to reach that node. The gain/loss 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.

The figure illustrates the discussion in Section 3.2. After a time 1 gain (point A), the investor takes a position such that, after a *poor* return from time 1 to time 2, she ends up with a time 2 gain slightly above zero (point B). After a time 1 loss (point C), she takes a position such that, after a *good* return from time 1 to time 2, she ends up with a time 2 gain that is also slightly above zero (point D'). The time 2 gain/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 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.

In our example, the disposition effect fails: the optimal time 1 allocation after a gain, $x_1(39.9)$, is 5.05, while the optimal time 1 allocation after a loss, $x_1(-24.3)$, is 3.06. Given that the investor starts with 4.0 shares at time 0, we see that she sells after a *loss*, rather than after a gain.

Figure 2 helps us understand *why* the disposition effect fails. For the investor to buy the stock at all at time 0, the stock must have a reasonably high expected return, so that $R_u - 1$ is somewhat larger than $1 - R_d$. This means two things. First, it 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, point A is further from the vertical axis than point C is. Second, when coupled with the mild concavity of $v(\cdot)$ in the region of gains, it means that the investor's optimal strategy, after a gain, is to gamble to the edge of the concave region; in other words, down to point B. However, it takes a larger allocation to gamble from point A to the edge of the concave region, as the investor's optimal strategy says she should, than it does to gamble from point C to the edge of the convex region, as her optimal strategy again says she should. The investor therefore takes more risk after a *gain* than after a loss, contrary to the disposition effect.

It is worth taking a moment to see where the traditional intuition linking prospect theory and the disposition effect can go wrong. The traditional intuition says that, 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. The investor therefore takes less risk after a gain than after a loss. In other words, she has a greater propensity to sell a stock after a gain than after a loss, and the disposition effect appears to hold.

The pitfall in this argument is the following. Since $v(\cdot)$ exhibits only *mild* concavity in the region of gains, the only reason an investor would take a small position in the risky asset after a gain is if the expected risky asset return is unattractive; in other words, if it is only slightly higher than the risk-free rate. In such a case, however, *the investor would not have*

bought the risky asset at time 0! For her to buy the risky asset 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 her almost to the edge of the concave region. This gamble is so large that the disposition effect fails: 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.[11]

This discussion also explains why, in Table 2, the disposition effect *does* sometimes hold; specifically, when there are many trading periods T and the expected risky asset return 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 gambles to the edge of the concave region. This relies on the fact that the expected risky asset return is quite high, which, in turn, is because otherwise, the investor would not buy the risky asset 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, lowering the investor's risk aversion. She is therefore willing to buy the risky asset 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 risky asset after a gain. As a result, the disposition effect can hold after all.

By identifying the situations in which the disposition effect does hold and the situations in which it does not, our analysis can help empirical researchers test our particular implementation of prospect theory. For example, since the disposition effect is more likely to hold when the expected return is low, we should empirically see more of a disposition effect among stocks with characteristics associated with low average returns. And given that the disposition effect is more likely to hold for a large number of trading periods T, we should see more of a disposition effect among traders who have a high T in mind when they buy stocks. Such traders could be identified, for example, by the high frequency of their trading.

4 Discussion

4.1 Robustness

In Sections 3.2 and 3.3, we found that, in a two-period setting, our implementation of prospect theory predicts the opposite of the disposition effect. How sensitive is this conclusion

---

[11]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 return on the risky asset has a binomial distribution. In Section 4.1, we discuss the robustness of our results to the use of a log-normal distribution.

to our modeling assumptions?

In our explanation for why the disposition effect fails, the first 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 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 investor is happy to buy the stock at time 0 because of its high expected return, but the potential time 1 gain is *smaller* in magnitude than the potential time 1 loss. The disposition effect may now hold: the allocation needed to gamble to the edge of the concave region after a small gain is lower than the allocation needed to gamble to the edge of the convex region after a 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).

Throughout Section 3, we set the parameters α and λ to the values estimated by Tversky and Kahneman (1992). Varying the degree of loss aversion λ has little effect on our results. Indeed, for the $(\mu, \sigma) = (1.1, 0.3)$ parameter values assumed in Section 3.3, and for $T = 2$ trading periods, we obtain identical optimal share allocations for any $\lambda \epsilon [1.25, 2.4]$. This is surprising at first: one might think that a lower value of λ would lead the investor to buy more shares at time 0, and yet, she does not. The reason is that, for most parameter values, the investor prefers her optimal gain/loss in the middle node at date 2 – point B in Figure 2 – to be slightly above zero. If, for λ below 2.25, she were to allocate more to the risky asset at time 0, that would mean a larger potential loss at time 1. After such a time 1 loss, the wealth constraint would prevent her from taking a position in the risky asset large enough to ensure that, even after a good return from time 1 to time 2, she ends up with a gain at time 2. She therefore maintains the same time 0 allocation.

When α is substantially lower than the benchmark level of 0.88, our results *do* change, and we do find a disposition effect. The parameter α governs the curvature of the value function $v(\cdot)$, both in the region of gains and in the region of losses. Lowering α increases concavity over gains and convexity over losses. Since, for lower α, the concavity over gains is higher, the investor takes a smaller position than before after a time 1 gain; she no longer gambles all the way to the edge of the concave region. Since the convexity over losses is now also higher, the investor takes a larger position than before after a time 1 loss. Once α falls sufficiently – in the case of $(\mu, \sigma) = (1.1, 0.3)$ and $T = 2$, once α falls below 0.77 – the investor holds more shares after a loss than after a gain, and we obtain a disposition effect. Figure 3, which has the same form as Figure 2, illustrates this result for $\alpha = 0.65$. After a time 1 loss, the investor takes a large position in the stock in order to gamble all the way to D'. After a time 1 gain, however, she doesn't need to take a very large position to gamble

down to B.

The model of Section 3 assumes that, once the investor has decided, at time 0, on the maximum amount W_0 she is willing to lose from trading the risky asset, she sticks to that decision. In results not reported here, we find that relaxing this assumption does not affect our conclusions. Specifically, suppose that, at time 1, the investor decides that she is willing to lose more than just the initial W_0, and that this decision is not anticipated at time 0. Since the wealth constraint no longer binds, the investor's allocation after a loss is given by the term before the comma in the "$\Delta W_1 < 0$" row of equation (42); in words, by the allocation that allows her to gamble to the edge of the convex region. The argument then proceeds as before: since the time 1 gain is larger in magnitude than the time 1 loss, it takes a larger position to gamble to the edge of the concave region than it does to gamble to the edge of the convex region. The disposition effect again fails to hold.

In our model, the expected stock return is constant over time. If instead, after a first-period gain, the investor for some reason lowers her estimate of the stock's expected return, she will be more inclined to sell, and we may see a disposition effect after all.

We note two things about this argument. First, 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. Second, and more important, if we are willing to assume that, after a gain, the investor lowers her estimate of the stock's expected return, then standard power utility preferences will deliver a disposition effect: we do not need to appeal to prospect theory at all! The goal of this paper is to examine the long-held view that prospect theory generates a disposition effect *without* auxiliary assumptions about changing beliefs. The analysis in Section 3 shows that this view can be formalized in some cases, but not in others.

Finally, for the sake of analytical tractability, our model assigns the risky asset return a binomial distribution, rather than a log-normal one. However, as $T \to \infty$, the binomial distribution converges to a log-normal one. When we analyze the investor's optimal trading strategy for large T, we find that our basic message remains intact: the disposition effect holds in some cases – specifically, for lower values of the expected risky asset return – but fails to hold once the expected return exceeds a certain level.

4.2   Related literature

To our knowledge, only two other papers attempt to formalize the link between prospect theory and the disposition effect. Hens and Vlcek (2005) study the portfolio problem of an investor with prospect theory preferences. Their analysis leads them to question, as do we, whether prospect theory predicts a disposition effect. At the same time, there are big

differences between our analysis and theirs. Hens and Vleck (2005) assume that there are only two trading periods; that the investor acts myopically, so that, at time 0, she does not take her time 1 decision into account; and that, each period, she invests either entirely in T-Bills or entirely in the risky asset. In this paper, we derive the trading behavior of a prospect theory investor in a far more general framework: one that allows for *any* number of trading periods and for full intertemporal optimization, and that places no restrictions on the set of feasible strategies; even in our two-period analysis, the investor takes her time 1 decision into account at time 0. Moreover, our ability to explore cases with many trading periods turns out to be important: qualitatively, the results are different for high T, in that the disposition effect tends to hold more often.

Gomes (2005) studies the two-period portfolio problem of an investor with preferences that are related to, but different from, prospect theory. Specifically, for losses below some specific point, he replaces the convex section of the prospect theory value function with a concave segment. He also sets α to 0.5, rather than to the 0.88 value estimated by Tversky and Kahneman (1992). Under these assumptions, his model predicts a disposition effect. Our paper shows that this result is special to the specific model he considers: for unmodified prospect theory, and for Tversky and Kahneman's (1992) parameters, the disposition effect always fails in two periods. As with Hens and Vlcek (2005), Gomes (2005) does not explore beyond the two-period case.

Another related paper is that of Kyle, Ou-Yang, and Xiong (2006). These authors 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, in that it does not take into account the investor's initial buying decision. As soon as we do, we recognize that the expected risky asset return must exceed a certain level. This, in turn, affects the frequency with which prospect theory can predict a disposition effect.

The results in Proposition 1 are related to those of Berkelaar, Kouwenberg, and Post (2004), who, in continuous time, solve the portfolio problem of a prospect theory investor with a horizon of T years. These authors make no mention at all of the disposition effect; rather, their focus is on how the investor's time 0 allocation varies with the length of the horizon T. We use a *discrete*-time framework because we want to be able to vary the frequency with which the investor can change her share holdings. This, in turn, allows us to study the way in which the link between prospect theory and the disposition effect depends on trading frequency.

4.3 Other applications

So far, we have applied our analysis to the trading of individual stocks. A number of recent papers uncover disposition-type evidence in other settings. Is our analysis also relevant in these other situations?

Genesove and Mayer (2001) find that homeowners are reluctant to sell their houses at prices below the original purchase price. They suggest that the concave-convex shape of the prospect theory value function may explain their evidence. The analysis in Section 3 shows that this explanation needs to be treated carefully: in its simplest implementation, prospect theory can sometimes make the opposite prediction, namely that people will be more reluctant to sell after a *gain* relative to purchase price than after a loss.[12]

Heath, Huddart, and Lang (1999) find that executives are more likely to exercise stock options when the underlying stock price exceeds a reference point – specifically, the stock's highest price over the previous year – than when it falls below that reference point. Does prospect theory predict these results, as the authors suggest? Not necessarily: our analysis shows that, in some cases, a prospect theory investor takes more risk after a *gain* than after a loss. However, there is an important caveat here. In our model, we know that the expected return on the risky asset is quite high: otherwise, the investor would not buy the asset in the first place. Executives, however, do not hold options because of an active decision that they made to buy them; they are simply endowed with them. Therefore, we cannot deduce, just from the fact that they hold options, that these options have a high expected return.

If the stock options have a *low* expected return, prospect theory can more easily predict the disposition-type effects in Heath et al. (1999). After a gain, the executive is in the concave region of the prospect theory value function. Even though the value function is only mildly concave in this region, the expected return on the options may be so low that the executive only wants to have a small position in them. After a loss, she is in the convex region, and so long as the expected return on the options exceeds the risk-free rate, she will at least gamble to the edge of the convex region. Overall, then, she may want to take less risk after a gain than after a loss, and may therefore exercise more options after a gain, consistent with Heath et al. (1999).

Our findings are also relevant to Coval and Shumway (2005), who 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

---

[12]One difference between buying a house and buying a stock is that a homebuyer is motivated not only by future investment performance, but also by factors such as the psychological utility of owning a home. So long as these other factors remain relatively constant over time, the models of Section 3 can still be used to predict propensity to sell; we simply need to adjust the risky asset's expected return upward to capture the positive effect that the additional factors have on the asset's desirability.

suggests that, in its simplest implementation, prospect theory can predict these results in some cases, but not in others. For example, we know from Section 3 that, if the expected return on the risky asset is high – in this context, if the expected return from trading is high – a trader will tend to take more risk after a *gain* than after a loss.

Recently, Grinblatt and Han (2005) and Frazzini (2006) have argued that, if a sizeable segment of the investor population exhibits the disposition effect, this could leave an imprint on asset prices and, in particular, could generate momentum and a post-earnings announcement drift in stock returns. Given that prospect theory has been closely associated with the disposition effect, Grinblatt and Han (2005) and Frazzini (2006) conjecture that, if momentum is driven by the disposition effect, then it may be ultimately by driven by the presence, in the economy, of investors with prospect theory preferences.

A full analysis of this conjecture is beyond the scope of this paper. Nevertheless, the partial equilibrium results of Section 3 suggest that, for our implementation of prospect theory, it may not hold in all cases. As we have seen, a prospect theory investor sometimes wants to take less risk after a *loss* than after a gain. In order to clear the market, then, expected stock returns may need to be higher after a loss than after a gain. This is the opposite of momentum.

5 Conclusion

One of the most striking portfolio puzzles is the "disposition effect": the tendency of individuals to sell stocks in their portfolios that have risen in value since purchase, rather than fallen in value. Perhaps the most prominent explanation for this puzzle is based on prospect theory. Despite its prominence, this explanation has received little formal scrutiny. We take up this task, and analyze the trading behavior of an investor with prospect theory preferences. We find that, at least for the simplest implementation of prospect theory, the link between these preferences and the disposition effect is not as obvious as previously thought: in some cases, prospect theory does indeed predict a disposition effect; but in other cases, it predicts the opposite.

We are keen to emphasize that, while our particular implementation of prospect theory does not always produce a disposition effect, it does do so in many cases. As such, it may turn out to be the right way of thinking about this particular pattern of trading. To determine this, future research could test some of the new predictions, discussed in Section 3, that come out of our model.

If the predictions of our model are not borne out in the data, we can turn our attention to other models, including, for example, alternative implementations of prospect theory. We

mentioned one such alternative in Section 2: a model in which the investor receives prospect theory utility from *realized* gains and losses. While preferences defined over realized gains and losses represent a significant departure from standard models, our analysis suggests that we may need this kind of structure in order to fully understand the disposition effect.

6   References

Barber, B. and T. Odean (2000), "Trading is Hazardous to Your Wealth: The Common Stock Performance of Individual Investors," *Journal of Finance* 55, 773-806.

Barberis, N., Huang M., and R. Thaler (2006), "Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing," forthcoming, *American Economic Review*.

Berkelaar, A., Kouwenberg, R., and T. Post (2004), "Optimal Portfolio Choice under Loss Aversion," *Review of Economics and Statistics* 86, 973-987.

Cox, J. and C. Huang (1989), "Optimal Consumption and Portfolio Policies when Asset Prices Follow a Diffusion Process," *Journal of Economic Theory* 39, 33-83.

Epstein, L. and S. Zin (1990), "First-order Risk Aversion and the Equity Premium Puzzle," *Journal of Monetary Economics* 26, 387-407.

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

Frazzini, A. (2006), "The Disposition Effect and Underreaction to News," *Journal of Finance*, forthcoming.

Genesove, D., and C. Mayer (2001), "Loss Aversion and Seller Behavior: Evidence from the Housing Market," *Quarterly Journal of Economics* 116, 1233-1260.

Gomes, F. (2005), "Portfolio Choice and Trading Volume with Loss Averse Investors," *Journal of Business* 78, 675-706.

Grinblatt, M. and B. Han (2005), "Prospect Theory, Mental Accounting, and the Disposition Effect," *Journal of Financial Economics* 78, 311-339.

Heath, C., Huddart, S., and M. Lang (1999), "Psychological Factors and Stock Option Exercise," *Quarterly Journal of Economics* 114, 601-627.

Hens, T. and M. Vlcek (2005), "Does Prospect Theory Explain the Disposition Effect?" Working paper, University of Zurich.

Kahneman, D. and A. Tversky (1979), "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* 47, 263-291.

Kyle, A., Ou-Yang, H., and W. Xiong (2006), "Prospect Theory and Liquidation Decisions," *Journal of Economic Theory*, forthcoming.

Odean, T. (1998), "Are Investors Reluctant to Realize Their Losses?" *Journal of Finance* 53, 1775-1798.

Rabin, M. (2000), "Risk Aversion and Expected Utility Theory: A Calibration Theorem," *Econometrica* 68, 1281-1292.

Shefrin, H. and M. Statman (1985), "The Disposition to Sell Winners too Early and Ride Losers too Long," *Journal of Finance* 40, 777-790.

Tversky, A. and D. Kahneman (1992), "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty* 5, 297-323.

Weber, M. and C. Camerer (1998), "The Disposition Effect in Securities Trading: An Experimental Analysis," *Journal of Economic Behavior and Organization* 33, 167-184.

7 Appendix

7.1 Proof of Proposition 1.

To prove the proposition, 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. Her final period wealth allocation could therefore be path dependent. To accommodate this possibility, we allow the investor to allocate wealth across stock price *paths*. As part of the proof of the proposition, we show that, in fact, her optimal wealth allocation is *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 $i \in \{1, 2, \cdots, M\}$. The ex-ante probability of path i is $\widehat{\pi}_i$, so that

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

The price of a contingent claim that pays \$1 at time T if the stock price evolves along path i is $\widehat{p}_i$. The state price density at the endpoint of the path is therefore

$$\widehat{q}_i = \frac{\widehat{p}_i}{\widehat{\pi}_i}.$$

In addition, W_0 is the investor's initial wealth at time 0, R_f is the per-period gross risk-free rate, and $\{\widehat{W}_i\}_{i=1}^{M}$ 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 $\widehat{q}_i = \widehat{q}_j$. We compute the state price density explicitly later in the proof.

Applying the reasoning of Cox and Huang (1989), we can rewrite problem (13)-(15) as

$$V = \max_{\{\widehat{W}_i\}} \sum_{i=1}^{M} \widehat{\pi}_i v(\widehat{W}_i - W_0 R_f^T), \tag{51}$$

subject to the budget constraint

$$\sum_{i=1}^{M} \widehat{\pi}_i \widehat{q}_i \widehat{W}_i = W_0 \tag{52}$$

and the non-negativity of wealth constraint

$$\widehat{W}_i \geq 0, \qquad 1 \leq i \leq M. \tag{53}$$

We write the reference wealth level $W_0 R_f^T$ as $\overline{W}$, for short, and define

$$\widehat{w}_i = \widehat{W}_i - \overline{W}$$

32

to be the investor's gain/loss relative to that reference level. The problem in (51)-(53) then becomes

$$V = \max_{\{\widehat{w}_i\}} \sum_{i=1}^{M} \widehat{\pi}_i v(\widehat{w}_i),$$ (54)

subject to

$$\sum_{i=1}^{M} \widehat{\pi}_i \widehat{q}_i \widehat{w}_i = 0$$ (55)

$$\widehat{w}_i \geq -\overline{W}, \qquad 1 \leq i \leq M.$$ (56)

We now prove the proposition through a series of lemmas.

Lemma 1. There exists at least one optimum.

Proof of Lemma 1. The set of feasible $\{\widehat{w}_i\}$, defined by constraints (55)-(56), is compact. The existence result then follows directly from Weierstrass' theorem.$\diamondsuit$

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

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

Lemma 2. If $\widehat{\pi}_M^{1-\alpha}\widehat{q}_M^{-\alpha} > \lambda\widehat{\pi}_1^{1-\alpha}\widehat{q}_1^{-\alpha}$, $\{\widehat{w}_i = 0\}_{i=1}^{M}$ cannot be the optimum.

Proof of Lemma 2. We prove the lemma by contradiction. Suppose that $\{\widehat{w}_i = 0\}_{i=1}^{M}$ is the optimum, so that the investor's value function takes the value $V = 0$. Consider the strategy

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

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

$$V' = \widehat{\pi}_M \frac{\widehat{\pi}_1^{\alpha}\widehat{q}_1^{\alpha}}{\widehat{\pi}_M^{\alpha}\widehat{q}_M^{\alpha}}x^{\alpha} - \lambda\widehat{\pi}_1 x^{\alpha}$$

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

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

Lemma 3. If the investor's optimal gain/loss $\widehat{w}_i$ 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 3. We prove the lemma by contradiction. Suppose that the investor's gain/loss is zero at the end of path i, so that $\widehat{w}_i = 0$. 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 any loss of

33

generality, that the investor's gain/loss is positive at the end of path j, so that $\widehat{w}_j = x > 0$. The contribution of paths i and j to total utility is

$$J = \widehat{\pi}_j x^\alpha.$$

We now modify this strategy by moving a small amount of wealth $\delta > 0$ from path j to path i, so that

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

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

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

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 $\widehat{w}_i = 0$ is optimal. Hence, the optimal gain/loss is different from zero at the end of all paths.$\diamondsuit$

Lemma 4. If the optimal allocation is non-zero, there must be at least one path at the end of which the gain/loss is $-\overline{W}$, so that the investor is wealth constrained.

Proof of Lemma 4. Suppose that $\vec{w} = (\widehat{w}_1, \widehat{w}_2, \cdots, \widehat{w}_M)$ is a non-zero optimal allocation, so that the value function $J(\vec{w}) > J(0) = 0$. Suppose also that the wealth constraint is never binding, so that $\widehat{w}_i > -\overline{W}$, $\forall i$. 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.$\diamondsuit$

In any non-trivial optimum – any optimum in which $\widehat{w}_i$ does not equal 0 for all i – there are three possible wealth allocations at the end of a path: a positive allocation ($\widehat{w}_i > 0$), an unconstrained negative allocation ($-\overline{W} < \widehat{w}_i < 0$), or a constrained negative allocation ($\widehat{w}_i = -\overline{W}$). In particular, from Lemma 3, we know that $\widehat{w}_i = 0$ cannot be an optimal allocation.

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

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

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 $\widehat{w}_i$ is

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

Since μ_i is associated with an inequality constraint,

$$\mu_i = 0 \ \text{ if } \ \widehat{w}_i > -\overline{W}, \qquad 1 \leq i \leq M,$$

and

$$\mu_i > 0 \ \text{ if } \ \widehat{w}_i = -\overline{W}, \qquad 1 \leq i \leq M.$$

34

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:

$$\widehat{w}_i = \left(\frac{\alpha}{\mu_0\widehat{q}_i}\right)^{\frac{1}{1-\alpha}} \tag{57}$$

2. An unconstrained negative wealth allocation:

$$\widehat{w}_i = -\left(\frac{\alpha\lambda}{\mu_0\widehat{q}_i}\right)^{\frac{1}{1-\alpha}}, \quad \text{if} \ \widehat{q}_i > \frac{\alpha\lambda}{\mu_0}\overline{W}^{-(1-\alpha)} \tag{58}$$

3. A constrained negative wealth allocation:

$$\widehat{w}_i = -\overline{W}, \quad \text{if} \ \widehat{q}_i > \frac{\alpha\lambda}{\mu_0}\overline{W}^{-(1-\alpha)}. \tag{59}$$

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

The multiplier μ_0 can be determined from the budget constraint:

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

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

$$\left(\frac{\alpha}{\mu_0}\right)^{\frac{1}{1-\alpha}} = \frac{\overline{W}\sum_{l=m+1}^{M}\widehat{\pi}_l\widehat{q}_l}{\sum_{i=1}^{k}\widehat{\pi}_i\widehat{q}_i^{-\frac{\alpha}{1-\alpha}} - \lambda^{\frac{1}{1-\alpha}}\sum_{j=k+1}^{m}\widehat{\pi}_j\widehat{q}_j^{-\frac{\alpha}{1-\alpha}}}. \tag{60}$$

The value function is therefore

$$\begin{aligned}
V &= \sum_{i=1}^{k}\widehat{\pi}_i\widehat{w}_i^{\alpha} - \sum_{j=k+1}^{m}\lambda\widehat{\pi}_j(-\widehat{w}_j)^{\alpha} - \sum_{l=m+1}^{M}\lambda\widehat{\pi}_l\overline{W}^{\alpha} \\
&= \overline{W}^{\alpha}\left[\left(\sum_{l=m+1}^{M}\widehat{\pi}_l\widehat{q}_l\right)^{\alpha}\left(\sum_{i=1}^{k}\widehat{\pi}_i\widehat{q}_i^{-\frac{\alpha}{1-\alpha}} - \lambda^{\frac{1}{1-\alpha}}\sum_{j=k+1}^{m}\widehat{\pi}_j\widehat{q}_j^{-\frac{\alpha}{1-\alpha}}\right)^{1-\alpha} - \lambda\sum_{l=m+1}^{M}\widehat{\pi}_l\right]. \tag{61}
\end{aligned}$$

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 5. It is not optimal to have a path with an unconstrained negative allocation, $\widehat{w}_i : -\overline{W} < \widehat{w}_i < 0$.

Proof of Lemma 5. From equation (61), 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.◇

Lemma 5 implies that we can write the value function as

$$V = \overline{W}^\alpha \left[\left(\sum_{i=1}^{k} \widehat{\pi}_i \widehat{q}_i^{-\frac{\alpha}{1-\alpha}} \right)^{1-\alpha} \left(\sum_{l=k+1}^{M} \widehat{\pi}_l \widehat{q}_l \right)^\alpha - \lambda \sum_{l=k+1}^{M} \widehat{\pi}_l \right]. \tag{62}$$

Lemma 6. Suppose that, in each period, a good stock return and a poor stock return are equally likely. In this case, 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 6. Given the equal probability of a good or poor return, each price path has the same probability:

$$\widehat{\pi}_i = \frac{1}{2^T}, \quad i = 1, \cdots, M.$$

The value function is then

$$V = 2^{-T} \overline{W}^\alpha \left[\left(\sum_{i=1}^{k} \widehat{q}_i^{-\frac{\alpha}{1-\alpha}} \right)^{1-\alpha} \left(\sum_{l=k+1}^{M} \widehat{q}_l \right)^\alpha - \lambda(M-k) \right]. \tag{63}$$

Suppose that $\widehat{q}_l < \widehat{q}_i$ for some $i \in \{1, \cdots, k\}$ and $l \in \{k+1, \ldots, M\}$. Equation (63) shows that assigning path l a positive wealth allocation and path i a constrained negative allocation $-\overline{W}$ 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.◇

Lemma 6 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 are not the same.

Another corollary of Lemma 6 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.

We now show that the final date wealth allocation cannot be path dependent at *any* final date node. We demonstrate this by showing that, if this is not true, then it is impossible to clear markets in an equilibrium model with both prospect theory agents and standard power utility agents.

Consider a Lucas exchange economy with $T+1$ dates, $t = 0, 1, \ldots, T$. There is a risky asset, which is a claim to a time T dividend of $D_0 G_1 \ldots G_T$. D_0 is known at time 0, while G_t is announced at time t. G_t is independently distributed for each t, and takes the value $D_u > 1$ with probability $\frac{1}{2}$ and $D_d < 1$ with probability $\frac{1}{2}$. News about the final dividend therefore evolves along a binomial tree.

There are two groups of investors: investors in the first group have the prospect theory objective function in (19); investors in the second group have a power utility function C^γ defined over final consumption. At date 0, the prospect theory investors hold a fraction κ of the initial endowment of the risky asset, and power utility investors hold the rest. We assume, as before, that markets are complete.

We denote prospect theory investors' optimal date T consumption at the end of path i as C_i^p. As discussed above, this consumption allocation depends on the path's state price density. We denote power utility investors' optimal date T consumption at the end of path i as C_i^n. A well-known property of power utility investors' optimal consumption is that it is strictly decreasing in a path's state price density.

Lemma 7. If paths i and j end at the same final date node, so that they have the same final date dividend, then they will have the same state price density and prospect theory investors will assign identical wealth allocations to the two paths.

Proof of Lemma 7. Suppose that prospect theory investors assign different wealth allocations to paths i and j. From our earlier discussion, one of the paths must have a constrained allocation of 0, while the other must have an allocation higher than $\overline{W}$. Without loss of generality, we let path i have the positive allocation and path j the zero allocation. Lemma 6 then implies $\widehat{q}_i \leq \widehat{q}_j$.

Market clearing implies that $C_i^n < C_j^n$. The optimization problem for the power utility investors then implies $\widehat{q}_i > \widehat{q}_j$. This gives a contradiction. Thus, the prospect theory investors must have the *same* wealth allocation for paths i and j. Market clearing means that the power utility investors also have the same wealth allocation for paths i and j. This implies that $\widehat{q}_i = \widehat{q}_j$. $\diamondsuit$

We have used a simple equilibrium argument to show that the wealth allocation at any date T node does not depend on the path by which the stock price arrives at that node. We can therefore revert to the notation of Section 3, where $P_{t,j}$, $W_{t,i}$, $x_{t,i}$, and $q_{t,i}$ denote the stock price, optimal wealth allocation, optimal share holding, 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, each period, $q_{t,i}$ either goes up by q_u or goes down by q_d. A standard property of the state price density is

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

which is equivalent to

$$\frac{1}{2}(q_u R_u + q_d R_d) = 1.$$

Another standard result is that the state price density satisfies a martingale property,

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

which is equivalent to

$$\frac{1}{2}(q_u + q_d) = 1.$$

We therefore obtain

$$q_u = \frac{2(1 - R_d)}{R_u - R_d}, \qquad q_d = \frac{2(R_u - 1)}{R_u - R_d}. \tag{64}$$

Since $(1 - R_d) < (R_u - 1)$, we have $q_u < q_d$.

We can now complete the proof. From (64), we know that the state price density increases as we go down the $T + 1$ date T nodes:

$$q_{T,1} < q_{T,2} < \ldots < q_{T,T+1}.$$

From Lemma 6, equation (57), and equation (60), 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 allocation of

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

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

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

and for the bottom $T + 1 - k^*$ nodes, she chooses an optimal wealth allocation of zero.

To determine k^*, we need to compute the investor's utility for each of the T possible values of k^*, $1 \leq k \leq T$, and to find the wealth allocation strategy that maximizes utility. Suppose that the investor chooses a positive wealth allocation for the top k nodes. From equation (62), utility is then given by

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

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

Given the optimal wealth allocations in the final period, $W_{T,j}$, we can compute optimal wealth allocations at all earlier dates using the state price density, as shown in equation (23).

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 stocks and B dollars of bonds. Her wealth at node $(t + 1, i)$ will therefore be $x_{t,i} P_{t+1,i} + B R_f$, and at node $(t + 1, i + 1)$, $x_{t,i} P_{t+1,i+1} + B R_f$. The difference must equal $W_{t+1,i} - W_{t+1,i+1}$, so that

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

which is equation (28). This completes the proof of the proposition.

7.2   Proof of Proposition 2.

Since the net risk-free rate is zero and the risky asset has a positive expected return ($R_u - 1 > 1 - R_d$), the investor will only take a long position in the risky asset. Her gain/loss at time 2 is

$$\Delta W_2 = \Delta W_1 + z_1(R_{1,2} - 1),$$

where $z_1 = x_1 P_1$ is the dollar amount she allocates to the risky asset at time 1. The investor is subject to a wealth constraint: her wealth cannot fall below zero after a poor return from time 1 to time 2, which means

$$
\begin{aligned}
W_2 &= W_1 + z_1(R_d - 1) \\
&= W_0 + \Delta W_1 + z_1(R_d - 1) \geq 0 \\
\Rightarrow \quad z_1 &\leq \frac{W_0 + \Delta W_1}{1 - R_d}.
\end{aligned}
$$

We analyze the investor's portfolio choice at time 1 depending on whether she has a positive, zero, or negative gain/loss at that time.

Case A: $\Delta W_1 < 0$. In this case, the investor will always have a loss after a negative shock at time 2. However, after a positive shock, she could have either a gain or a loss, depending on the size of z_1. We consider two scenarios.

Scenario I: The investor takes a position small enough so that ΔW_2 is always negative ($z_1 < \frac{-\Delta W_1}{R_u - 1}$). Her expected utility is then

$$L_1(z_1) = -\frac{\lambda}{2}[-\Delta W_1 - (R_u - 1)z_1]^\alpha - \frac{\lambda}{2}[-\Delta W_1 + (1 - R_d)z_1]^\alpha, \quad \forall\, z_1 \in \left[0, \frac{-\Delta W_1}{R_u - 1}\right].$$

Differentiating, we obtain

$$L_1'(z_1) = \frac{\alpha\lambda}{2}(R_u - 1)[-\Delta W_1 - (R_u - 1)z_1]^{-(1-\alpha)} - \frac{\alpha\lambda}{2}(1 - R_d)[-\Delta W_1 + (1 - R_d)z_1]^{-(1-\alpha)},$$

so that

$$L_1'(z_1) > 0 \quad \forall\ z_1 \in \left[0, \frac{-\Delta W_1}{R_u - 1}\right].$$

$L_1(z_1)$ is therefore monotonically increasing in this region.

Scenario II: The investor takes a large enough position at time 1 so that ΔW_2 is positive after a positive shock at time 2 ($z_1 > \frac{-\Delta W_1}{R_u - 1}$). Her expected utility is then

$$L_1(z_1) = \frac{1}{2}[\Delta W_1 + (R_u - 1)z_1]^\alpha - \frac{\lambda}{2}[-\Delta W_1 + (1 - R_d)z_1]^\alpha, \quad \forall\, z_1 > \frac{-\Delta W_1}{R_u - 1}.$$

Differentiating, we obtain

$$L_1'(z_1) = \frac{\alpha}{2}(R_u - 1)[\Delta W_1 + (R_u - 1)z_1]^{-(1-\alpha)} - \frac{\alpha\lambda}{2}(1 - R_d)[-\Delta W_1 + (1 - R_d)z_1]^{-(1-\alpha)}.$$

It is easy to verify that $L'_1\left(\frac{-\Delta W_1}{R_u-1}\right) > 0$. When $g > \lambda^{\frac{1}{\alpha}}$, where g is defined in (41), $L'_1(z_1) = 0$ does not have a solution in the range $\left(\frac{-\Delta W_1}{R_u-1}, \infty\right)$. $L_1(z_1)$ is therefore monotonically increasing for positive z_1 and the investor takes the largest position allowed by the wealth constraint, $z_1 = \frac{W_0+\Delta W_1}{1-R_d}$. When $g < \lambda^{\frac{1}{\alpha}}$, there is a unique solution z^* to $L'_1(z_1) = 0$ for $z_1 > \frac{-\Delta W_1}{R_u-1}$, namely

$$z^* = \left(\frac{R_u - R_d}{\left(\frac{R_u-1}{\lambda(1-R_d)}\right)^{\frac{1}{1-\alpha}} + 1} - (1 - R_d)\right)^{-1} (-\Delta W_1).$$

Thus, $L_1(z_1)$ reaches its maximum at z^*. The optimal share holding in this case is given by

$$z_1(\Delta W_1) = \min\left(z^*, \frac{W_0 + \Delta W_1}{1 - R_d}\right).$$

Case B: $\Delta W_1 = 0$. In this case, the investor will always have a gain after a positive time 2 shock and a loss after a negative time 2 shock. Thus, her expected utility is

$$L_2(z_1) = \frac{1}{2}[(R_u - 1)^\alpha - \lambda(1 - R_d)^\alpha]z_1^\alpha, \quad \forall z_1 > 0.$$

The investor therefore takes a zero position if $g < \lambda^{1/\alpha}$; otherwise, she takes the largest position allowed by the wealth constraint, $z_1 = \frac{W_0}{1-R_d}$.

Case C: $\Delta W_1 > 0$. In this case, the investor will always have a gain after a positive shock at time 2. However, after a negative shock, she could have either a gain or a loss, depending on the size of z_1.

Scenario I: The investor takes a position small enough so that ΔW_2 is always positive, $\left(z_1 < \frac{\Delta W_1}{1-R_d}\right)$. Her expected utility is then

$$L_3(z_1) = \frac{1}{2}[\Delta W_1 + (R_u - 1)z_1]^\alpha + \frac{1}{2}[\Delta W_1 - (1 - R_d)z_1]^\alpha, \quad \forall z_1 \in \left[0, \frac{\Delta W_1}{1 - R_d}\right].$$

Differentiating, we obtain

$$L'_3(z_1) = \frac{\alpha}{2}(R_u - 1)[\Delta W_1 + (R_u - 1)z_1]^{-(1-\alpha)} - \frac{\alpha}{2}(1 - R_d)[\Delta W_1 - (1 - R_d)z_1]^{-(1-\alpha)}.$$

It is easy to verify that $L'_3(0) > 0$ and that there is a unique solution $z^{**} \in \left[0, \frac{\Delta W_1}{1-R_d}\right]$ to $L'_3(z_1) = 0$, which is

$$z^{**} = \left((1 - R_d) + \frac{R_u - R_d}{\left(\frac{R_u-1}{1-R_d}\right)^{\frac{1}{1-\alpha}} - 1}\right)^{-1} \Delta W_1.$$

Therefore, inside the region $\left[0, \frac{\Delta W_1}{1-R_d}\right]$, $L_3(z_1)$ reaches its maximum at z^{**}. Direct substitution gives

$$L_3(z^{**}) = \frac{(\Delta W_1)^\alpha}{2}\left(\frac{1+g}{1+g^{\frac{\alpha}{1-\alpha}}}\right)^\alpha \left(g^{\frac{\alpha^2}{1-\alpha}} + g^{-\alpha}\right).$$

40

Scenario II: The investor takes a large enough position so that ΔW_2 is negative after a negative shock at time 2 ($z_1 > \frac{\Delta W_1}{1-R_d}$). Her expected utility is then

$$L_3(z_1) = \frac{1}{2}[\Delta W_1 + (R_u - 1)z_1]^\alpha - \frac{\lambda}{2}[-\Delta W_1 + (1 - R_d)z_1]^\alpha, \quad \forall \; z_1 > \frac{\Delta W_1}{1 - R_d}.$$

Differentiating, we obtain

$$L_3'(z_1) = \frac{\alpha}{2}\left\{ (R_u - 1)^\alpha \left[z_1 + \frac{\Delta W_1}{(R_u - 1)}\right]^{-(1-\alpha)} - \lambda(1 - R_d)^\alpha \left[z_1 - \frac{\Delta W_1}{(1 - R_d)}\right]^{-(1-\alpha)} \right\}.$$

$L_3'(z_1)$ starts with a negative value at $\frac{\Delta W_1}{1-R_d}$. When $g < \lambda^{\frac{1}{\alpha}}$, $L_3'(z_1) = 0$ does not have a solution in the range $(\frac{\Delta W_1}{1-R_d}, \infty)$. $L_3(z_1)$ is therefore monotonically decreasing in this range, and the investor's optimal position is given by z^{**}. When $g > \lambda^{\frac{1}{\alpha}}$, $L_3'(z_1)$ starts out negative in the range $(\frac{\Delta W_1}{1-R_d}, \infty)$, but eventually turns positive as z_1 becomes large. Thus, it is possible that $L_3(z_1)$ in this region can become higher than $L_3(z^{**})$. If this occurs, the investor takes the largest position allowed by the wealth constraint, $z_1 = \frac{W_0+\Delta W_1}{1-R_d}$. In this situation, the investor's expected utility is

$$L_3\left(\frac{W_0 + \Delta W_1}{1 - R_d}\right) = \frac{1}{2}W_0^\alpha \left\{ \left[(g + 1)\frac{\Delta W_1}{W_0} + g\right]^\alpha - \lambda \right\}.$$

Thus, if

$$h(\Delta W_1) = \left[(g + 1)\frac{\Delta W_1}{W_0} + g\right]^\alpha - \left(\frac{1 + g}{1 + g^{\frac{\alpha}{1-\alpha}}}\right)^\alpha \left(g^{\frac{\alpha^2}{1-\alpha}} + g^{-\alpha}\right)\left(\frac{\Delta W_1}{W_0}\right)^\alpha - \lambda > 0,$$

the investor takes the maximum position allowed by the wealth constraint, and otherwise takes a position z^{**}. It is straightforward to check that $h(0) > 0$, that $h(\cdot)$ is monotonically decreasing for $\Delta W_1 > 0$, and that $h(\Delta W_1) = 0$ has a unique positive solution $\Delta W_1 = \Delta W^*$. The investor therefore takes a position z^{**} for $\Delta W_1 > \Delta W^*$, and otherwise takes the maximum position allowed by the wealth constraint.

Table 1: We solve a portfolio problem with a risk-free asset and a risky asset, and five dates, $t = 0, \ldots, 4$. The investor has prospect theory preferences defined over her accumulated trading profit at 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 in the risky asset and the optimal wealth, respectively. The net risk-free rate is zero and the annual net expected return and standard deviation of the risky asset are 0.1 and 0.3, respectively.

		$P_{t,i}$						$q_{t,i}$		
				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

		$x_{t,i}$						$W_{t,i}$		
				-						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

42

Table 2: For a given (μ, T) pair, we construct an artificial dataset of how 10,000 investors with prospect theory preferences, each of whom owns N_S stocks, each of which has an annual gross expected return μ, would trade those stocks over T periods. For each (μ, T) pair, we use the artifical dataset to compute PGR and PLR, where PGR is the proportion of gains realized by all investors over the entire trading period, and PLR is the proportion of losses realized. The table reports "PGR/PLR" for each (μ, T) pair. Boldface type identifies cases where the disposition effect fails (PGR $<$ PLR). A hyphen indicates that the expected return is so low that the investor does not buy any stock at all.

μ	$T = 2$	$T = 4$	$T = 6$	$T = 12$
1.03	-	-	-	.55/.50
1.04	-	-	.54/.52	.54/.52
1.05	-	-	.54/.52	.59/.45
1.06	-	.70/.25	.54/.52	.58/.47
1.07	-	.70/.25	.54/.52	.57/.49
1.08	-	.70/.25	**.48/.58**	**.47/.60**
1.09	-	**.43/.70**	**.48/.58**	**.46/.61**
1.10	**0.0/1.0**	**.43/.70**	**.48/.58**	**.36/.69**
1.11	**0.0/1.0**	**.43/.70**	**.49/.58**	**.37/.68**
1.12	**0.0/1.0**	**.28/.77**	**.23/.81**	**.40/.66**
1.13	**0.0/1.0**	**.28/.77**	**.24/.83**	**.25/.78**

Table 3: We solve a two-period portfolio problem with a risk-free asset and a risky asset. The investor has prospect theory preferences defined over her accumulated trading profit at the final date. $x_{0,1}$ is the optimal share allocation at time 0; $x_{1,1}$ and $x_{1,2}$ are the optimal time 1 share allocations after a gain and after a loss, respectively. μ is the annual gross expected return on the risky asset.

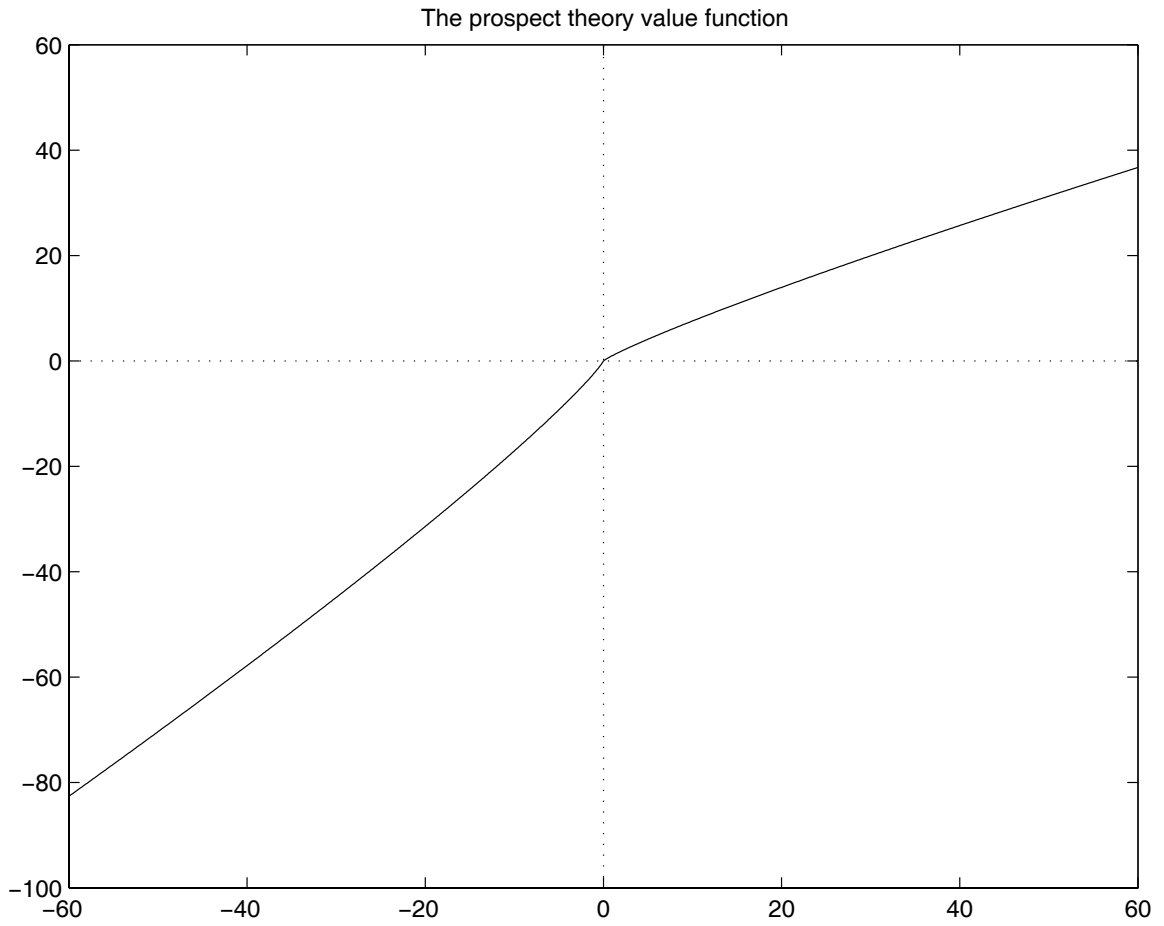| μ	$x_{0,1}$	$x_{1,1}$	$x_{1,2}$
1.06	0.00	0.00	0.00
1.08	0.00	0.00	0.00
1.10	4.00	5.05	3.06
1.12	4.55	6.51	2.99
1.14	5.17	8.30	2.94
1.16	5.88	10.60	2.92
1.18	6.72	13.66	2.90
1.20	7.73	17.85	2.88

Figure 1. The graph shows the form of the prospect theory value function $v(\cdot)$ proposed by Tversky and Kahneman (1992).

45

Time 1 and time 2 gains/losses plotted on the value function

Figure 2. The graph plots the gains and losses experienced by an investor with prospect theory preferences against her final period utility function. Point A is the potential time 1 gain and points B and B' are the time 2 gains that her corresponding time 1 allocation could lead to. Point C is the potential time 1 loss and points D and D' are the time 2 gains/losses that her corresponding time 1 allocation could lead to.

Time 1 and time 2 gains/losses plotted on the value function

Figure 3. The graph plots the gains and losses experienced by an investor with prospect theory preferences against her final period utility function. Point A is the potential time 1 gain and points B and B' are the time 2 gains that her corresponding time 1 allocation could lead to. Point C is the potential time 1 loss and points D and D' are the time 2 gains/losses that her corresponding time 1 allocation could lead to. The parameter α is set to 0.65, lower than the benchmark value of 0.88, thereby increasing the concavity (convexity) of the utility function in the region of gains (losses).