The Loss Aversion / Narrow Framing Approach to the Equity Premium Puzzle

Nicholas Barberis and Ming Huang
Yale University and Cornell University

October 2005*

Abstract

We review one recent approach to the equity premium puzzle. The key elements of this approach are loss aversion and narrow framing, two well-known features of decision-making under risk in experimental settings. In equilibrium, models that incorporate these ideas can generate a large equity premium and a low and stable risk-free rate, even when consumption growth is smooth and only weakly correlated with the stock market. Moreover, they can do so for parameter values that correspond to sensible predictions about attitudes to independent monetary gambles. The analysis for the equity premium also has implications for a closely related portfolio puzzle, the stock market participation puzzle. We suggest some possible directions for future research.

*This essay is being prepared for the *Handbook of Investments: Equity Premium,* edited by Rajnish Mehra. We thank Raj Mehra for inviting us to contribute a chapter, and are grateful to him for valuable feedback. Comments are welcome at nick.barberis@yale.edu and mh375@cornell.edu.

1 Introduction

One of the best-known stock market puzzles is the equity premium puzzle, which asks why investors historically appear to have demanded a high average return on stocks, relative to T-Bills (Mehra and Prescott, 1985). In this essay, we discuss one recent approach to addressing this puzzle. The broad theme of this approach is that we may be able to improve our understanding of how people evaluate stock market risk by looking at how they evaluate risk in experimental settings. More specifically, this approach argues that *loss aversion* and *narrow framing*, two of the most important ideas to emerge from the experimental literature on decision-making under risk, may also play an important role in the stock market setting.

Loss aversion is a central feature of Kahneman and Tversky's (1979) prospect theory – a descriptive theory, based on extensive experimental evidence, of how people evaluate risk. In this theory, the carriers of value are not absolute levels of wealth, but rather, gains and losses measured relative to a reference point. Loss aversion is the specific finding that people are much more sensitive to losses – even small losses – than to gains of the same magnitude.

To understand narrow framing, recall that under traditional utility functions defined over consumption or total wealth, the agent evaluates a new gamble by first mixing it with the other risks he is already facing and then checking whether the combination is attractive. Narrow framing, by contrast, is the phenomenon documented in experimental settings whereby, when people are offered a new gamble, they evaluate it in isolation, separately from their other risks. In other words, they act as if they get utility *directly* from the outcome of the gamble, even if the gamble is just one of many that determine their overall wealth risk. This contrasts with traditional specifications, in which the agent would only get utility from the outcome of the gamble *indirectly*, via its contribution to his total wealth.

Motivated by these ideas, some recent papers propose that people are *loss averse over changes in the value of their stock market holdings*. In other words, even if stock market risk is just one of many risks that determine their overall wealth risk – others being labor income risk and housing risk, say – they still get utility directly from stock market fluctuations (narrow framing) and are more sensitive to losses than to gains (loss aversion). For reasons we discuss below, most implementations also assume that people focus on *annual* gains and losses. Very informally, then, people evaluate stock market risk by saying: "well, stocks could go up over the next year, with roughly 50% probability; but they could also go down, with roughly the same probability. I'm much more sensitive to losses than to gains, so this doesn't look like an attractive risk to me." According to the approach we describe in this essay, it is this sort of thinking that leads the investing population to charge a high premium for holding the market supply of stocks; in other words, to charge a high equity premium.

Why should financial economists be interested in this particular approach to the equity premium puzzle? What are its selling points? In this survey, we emphasize two. First, the framework we describe can generate a high equity premium while also matching *other* aspects of the data, such as the low and stable risk-free rate, the low volatility of consumption growth and the low correlation of stock returns and consumption growth. With some additional structure, it can also match the high volatility and time-series predictability of stock returns.

A second benefit of our approach is that it can address the equity premium puzzle for preference parameters that are "reasonable," by which we mean parameters that also make sensible predictions about attitudes to independent monetary gambles. This is important because it was, in part, the difficulty researchers encountered in reconciling the high average return on stocks with reasonable attitudes to large-scale monetary gambles that launched the equity premium literature in the first place.

The approach we describe here was first proposed by Benartzi and Thaler (1995). In their framework, the investor is loss averse over fluctuations in the value of his financial wealth, which, since financial wealth is just one component of total wealth, constitutes narrow framing. One drawback of this framework is that, since the investor gets no direct utility at all from consumption or total wealth, consumption plays no role, making it hard to check how well the model describes the joint properties of stock returns and consumption growth.

Benartzi and Thaler's work therefore opens up a new challenge: to build and evaluate more realistic models in which, even if the investor gets utility from fluctuations in the value of one component of his wealth, he also gets some utility from consumption. In large part, this essay surveys the progress that has been made on this front, drawing primarily on Barberis, Huang and Santos (2001), Barberis, Huang and Thaler (2003), and Barberis and Huang (2004).[1]

The story we tell in this paper is a simple one: investors require a high equity premium because any drop in the stock market over the next year will bring them direct disutility. To some readers, this story may be *too* simple, in that the distance between assumption and conclusion may appear too close for comfort. We understand this point of view, and we agree that in order for the loss aversion / narrow framing framework to gain more currency in the profession, its unique predictions must be tested and confirmed. Fortunately, tests of our approach *are* starting to appear in the literature, and we discuss some of them at the end of the essay. Even before the outcome of these tests is known, however, there is a contribution in this paper that even a skeptical reader can appreciate, namely a methodological one: we show how loss aversion and narrow framing can be incorporated into more traditional models of asset pricing and we use the new models to better understand the predictions of our approach.

[1] Other investigations of loss aversion and narrow framing in a financial context include Berkelaar, Kouwenberg and Post (2004) and Gomes (2005).

In Section 2, we discuss loss aversion and narrow framing in more detail, examining both the evidence they are derived from and some of the interpretations they are given. In Section 3, we show that, once embedded into more traditional utility functions, these features can generate a high equity premium and a low and stable risk-free rate, even when consumption growth is smooth and only weakly correlated with stock returns; and moreover, that they can do so for parameter values that also make sensible predictions about attitudes to both large-scale and small-scale monetary gambles. We highlight the crucial role that narrow framing plays in our results by showing that without this feature, the approach loses many of its advantages. In Section 4, we show that our analysis of the equity premium also has implications for a portfolio choice puzzle, the stock market participation puzzle. Section 5 considers various extensions of the basic framework, while Section 6 concludes and discusses possible directions for future research.

Since loss aversion and narrow framing are the defining features of the approach we describe here, the framework should, strictly speaking, be called the "loss aversion and narrow framing" approach to the equity premium puzzle. Given that narrow framing is the more distinctive of the two ingredients, we sometimes abbreviate this to the "narrow framing" approach.[2]

2 Loss Aversion and Narrow Framing

Loss aversion is a central feature of Kahneman and Tversky's (1979) prospect theory, a descriptive theory of decision-making under risk, in which the carriers of value are not absolute wealth levels, but rather, gains and losses measured relative to a reference point. Loss aversion is a greater sensitivity to losses – even small losses – than to gains of the same magnitude, and is represented by a kink in the utility function.

The most basic evidence for loss aversion is the fact that people tend to reject gambles of the form

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

to be read as "win \$110 with probability $\frac{1}{2}$, lose \$100 with probability $\frac{1}{2}$, independent of other risks" (Kahneman and Tversky, 1979, Tversky and Kahneman, 1992). It is hard to explain this evidence with differentiable utility functions, whether expected utility or non-expected utility, because the very high local risk aversion required to do so typically predicts an implausibly high level of aversion to *large*-scale gambles (Epstein and Zin, 1990, Rabin,

[2]Benartzi and Thaler (1995) use the label "myopic loss aversion". By using this phrase, they emphasize the investor's sensitivity to losses (loss aversion) and his focus on annual changes (myopia), but not the narrow framing. As we will see, narrow framing is more crucial to our results than the annual evaluation of gains and losses, and so we prefer to emphasize the narrow framing feature while playing down the myopia.

2000, Barberis, Huang and Thaler, 2003).[3]

The classic demonstration of narrow framing is due to Tversky and Kahneman (1981), who ask 150 subjects the following question:

Imagine that you face the following pair of concurrent decisions. First examine both decisions, then indicate the options you prefer:

Choice (I) Choose between:

A. a sure gain of $240

B. 25% chance to gain $1,000 and 75% chance to gain nothing

Choice (II) Choose between:

C. a sure loss of $750

D. 75% chance to lose $1,000 and 25% chance to lose nothing.

Tversky and Kahneman (1981) report that 84% of subjects chose A, with only 16% choosing B, and that 87% chose D, with only 13% choosing C. In particular, 73% of subjects chose the combination A&D, namely

$$\text{25\% chance to win \$240,} \quad \text{75\% chance to lose \$760,} \tag{2}$$

which is surprising, given that this choice is dominated by the combination B&C, namely

$$\text{25\% chance to win \$250,} \quad \text{75\% chance to lose \$750.} \tag{3}$$

It appears that instead of focussing on the *combined* outcome of decisions I and II – in other words, on the outcome that determines their final wealth – subjects are focussing on the outcome of each decision separately. Indeed, subjects who are asked *only* about decision I do overwhelmingly choose A; and subjects asked *only* about decision II do overwhelmingly choose D.

In more formal terms, it appears that we cannot model the typical subject as maximizing a utility function defined only over total wealth. Rather, his utility function appears to depend *directly* on the outcome of each of decisions I and II, rather than just indirectly, via

[3]There is also strong evidence of what Thaler (1980) calls an "endowment effect," which can be thought of as loss aversion in the absence of uncertainty. Kahneman, Knetsch and Thaler (1990) conduct a series of experiments in which subjects are either given some object such as a coffee mug and then asked if they would be willing to sell it, or not given the mug and then offered the chance to buy one. The authors find that mug owners demand more than twice as much to sell their mugs as non-owners are willing to pay to acquire one.

the contribution of each decision to overall wealth. As such, this is an example of narrow framing.

More recently, Barberis, Huang and Thaler (2003) have argued that the commonly observed rejection of the gamble in (1) is not only evidence of loss aversion, but of narrow framing as well. To see why loss aversion on its own cannot explain this behavior, note that most of the subjects who are offered this gamble are typically already facing *other* kinds of risk – labor income risk, housing risk, or financial market risk, say. In the absence of narrow framing, they must therefore evaluate the 110/100 gamble by mixing it with these other risks and then checking if the combination is attractive. It turns out that the combination *is* almost always attractive: since the 110/100 gamble is independent of other risks, it offers useful diversification benefits, which, even if loss averse, people can enjoy. The rejection of the 110/100 gamble therefore strongly suggests that people are *not* fully merging it with their other risks, but that to some extent, they are evaluating it in isolation; in other words, that they are framing it narrowly.

By the same token, *any* evidence of aversion to a small, independent, actuarially favorable risk is likely to stem from narrow framing. Examples of such evidence in the field are the high premia consumers pay for telephone wiring insurance and the low deductibles chosen in automobile insurance contracts (Cicchetti and Dubin, 1994, Rabin and Thaler, 2001, Grgeta and Thaler, 2003). Similar evidence in experimental data can be found in Bossaerts, Plott and Zame (2003).[4]

Motivated by these ideas, some recent papers propose that people are *loss averse over changes in the value of their stock market holdings*. In other words, even if stock market risk is just one of many risks that determine their overall wealth risk – others being labor income risk and housing risk, say – they still get utility directly from stock market fluctuations (narrow framing) and are more sensitive to losses than to gains (loss aversion).

Is it plausible that people might frame stock market risk narrowly? To answer this, it may be helpful to think about how narrow framing can be interpreted. One interpretation is that it stems from non-consumption utility, such as regret. Regret is the pain we feel when we realize that we would be better off today if we had taken a different action in the past. Even if a gamble that an agent accepts is just one of many risks that he faces, it is still linked to a specific decision, namely the decision to accept the gamble. As a result, it exposes the agent to possible future regret: if the gamble turns out badly, he may regret the decision to accept it. Consideration of non-consumption utility therefore leads quite naturally to preferences that depend *directly* on the outcomes of specific gambles the agent faces.

[4]For more evidence of narrow framing, see Kahneman and Tversky (1983), Tversky and Kahneman (1986), Redelmeier and Tversky (1992), Kahneman and Lovallo (1993) and Read, Lowenstein and Rabin (1999).

A second interpretation of narrow framing is proposed by Kahneman (2003). He argues that many decisions are made intuitively rather than through effortful reasoning. Since intuitive thoughts are by nature spontaneous, they are heavily shaped by the features of the situation at hand that come to mind most easily; to use the technical term, by the features that are most "accessible." In this view, narrow framing occurs because sometimes, when an agent evaluates a new gamble, the distribution of the gamble, considered separately, is much more accessible than the distribution of his overall wealth once the new gamble has been merged with his other risks. The fact that the distribution of the gamble, taken alone, is so accessible, means that that distribution plays a more important role in decision-making than would be predicted by traditional utility functions defined only over wealth or consumption.

In Tversky and Kahneman's (1981) example, the outcome of each one of choices A, B, C or D is highly accessible. Much less accessible, though, is the *overall* outcome once two choices – A&D, say, or B&C – are combined: the distributions in (2) and (3) are less "obvious" than the distributions of A, B, C and D given in the original question. As a result, the outcome of each of decisions I and II plays a bigger role in decision-making than predicted by traditional utility functions. Similar reasoning can be applied in the case of the 110/100 gamble.

It seems to us that both the "regret" and "accessibility" interpretations of narrow framing *do* apply naturally to the stock market. Allocating some fraction of his wealth to the stock market constitutes a specific action on the part of the agent – one that he may later regret if his stock market gamble turns out poorly.[5] Alternatively, given our daily exposure, through newspapers, books and other media, to large amounts of information about the distribution of the stock market, such information is very accessible. Much less accessible is any information as to the distribution of future outcomes once stock risk is merged with the other kinds of risk that people face. Judgments about how much to invest in stocks might therefore be made, at least in part, using a narrow frame.

The accessibility interpretation of narrow framing also provides a rationale for why investors might focus on *annual* gains and losses in the stock market. Much of the public discussion about the historical performance of different asset classes is couched in terms of annual returns, making the annual return distribution particularly accessible.[6]

[5]Of course, investing in T-Bills may also lead to regret if the stock market goes up in the meantime. Regret is typically thought to be stronger, however, when it stems from having taken an action – for example, actively moving one's savings from the default option of a riskless bank account to the stock market – than from having not taken an action – for example, leaving one's savings in place at the bank. In short, errors of commission are more painful than errors of omission.

[6]Clever tests of this logic can be found in Gneezy and Potters (1997) and Thaler et al. (1997). The latter paper, for example, asks subjects how they would allocate between a risk-free asset and a risky asset over a long time horizon – 30 years, say. The key manipulation is that some subjects are given the distribution of asset returns over short horizons – monthly returns, say – while others are given a long-term return distribution – the distribution of 30-year returns, say. Since they have the same decision problem, the two

While Tversky and Kahneman's (1981) experiment provides conclusive evidence of narrow framing, it is also somewhat extreme, in that in this example, narrow framing leads subjects to choose a dominated alternative. In more general situations, this will not be the case. All the same, Tversky and Kahneman's (1981) example does raise the concern that when applied to asset pricing, narrow framing might give rise to arbitrage opportunities. To avoid this problem, we focus on applications to *absolute* pricing – in other words, to the pricing of assets, like the aggregate stock market, that lack perfect substitutes – because in such situations, there are no riskless arbitrage opportunities. We would not expect narrow framing to have much useful application to relative pricing: in this case, any impact that narrow framing had on prices would create an arbitrage opportunity that could be quickly exploited.

The normative status of narrow framing depends on the interpretation it is given. If it stems from non-consumption utility, such as regret, then it *can* be normatively acceptable: the agent simply gets utility from things other than consumption, and takes this into account when making decisions. Since he is acting optimally, there is no reason to expect his behavior to change over time. Narrow framing is therefore likely to be a permanent feature of preferences and if it leads the agent to charge a high equity premium today, then it will lead him to charge a high equity premium in the future as well.

If, however, narrow framing stems from intuitive thinking and from basing decisions only on "accessible" information, it becomes less acceptable from a normative standpoint. Given his preferences, the agent would be happier with a different decision rule, but has failed to go through the effortful reasoning required to uncover that rule. In this case, we would expect the agent's behavior to change over time, as he learns that his intuitive thinking is leading him astray, and either through his own efforts, or by observing the actions of others, discovers a better decision rule. If accessibility-based narrow framing is driving the equity premium, we would expect the premium to fall over time as investors gradually switch away from narrow framing.

Our discussion has treated loss aversion and narrow framing as two distinct phenomena. Recent work, however, suggests that they may form a natural pair, because in those situations where people exhibit loss aversion, they often also exhibit narrow framing. For example, as argued above, the rejection of the 110/100 gamble in (1) is evidence not only of loss aversion, but of narrow framing as well.

---

groups of subjects should make similar allocation decisions: those subjects given the shorter-term return distribution should simply use it to infer the more directly relevant longer-term distribution. In fact, these subjects allocate substantially less to the risky asset, suggesting that they are simply falling back on the distribution that is most accessible to them, namely the short-term return distribution they were given. Since losses occur more often in high frequency data, they perceive the risky asset to be especially risky and allocate less to it.

Kahneman (2003) suggests an explanation for why loss aversion and narrow framing might appear in combination like this. He argues that prospect theory captures the way people act when making decisions intuitively, rather than through effortful reasoning. Since narrow framing is also thought to derive, at least in part, from intuitive decision-making, it is natural that prospect theory, and therefore also loss aversion, would be used in parallel with narrow framing.

3 The Equity Premium

In this section, we discuss various ways of modeling loss aversion and narrow framing, and then demonstrate the advantages, from the perspective of addressing the equity premium puzzle, of a model in which investors frame stock market risk narrowly. Specifically, in Section 3.2, we show that such a model can generate a high equity premium at the same time as a low and stable risk-free rate, even when consumption growth is smooth and only weakly correlated with stock returns; and then, in Section 3.3, that it can do so while also making reasonable predictions about attitudes to large-scale monetary gambles.

3.1 Modeling loss aversion and narrow framing

Benartzi and Thaler (1995) are the first to apply loss aversion and narrow framing in the context of the aggregate stock market. They consider an investor who is loss averse over changes in the value of his financial wealth, defined here as holdings of T-Bills and stocks. Since financial wealth is just one component of overall wealth – others being human capital and housing wealth – defining utility directly over fluctuations in financial wealth constitutes narrow framing.

Benartzi and Thaler (1995) argue that, in equilibrium, their investor will charge a high equity premium. In simple terms, the high volatility of stock returns leads to substantial volatility in returns on financial wealth. Given that he is more sensitive to losses than to gains, these fluctuations in his financial wealth cause the investor substantial discomfort. As a result, he will only hold the market supply of stocks if compensated by a high average return.

A weakness of Benartzi and Thaler's (1995) framework is that, since the investor gets direct utility *only* from changes in the value of his financial wealth, and none at all from consumption or total wealth, consumption plays no role, making it hard to check how well the model describes the joint properties of stock returns and consumption growth. An important challenge therefore remains: to build and evaluate a more realistic model in which, even if the investor gets utility from fluctuations in the value of one component of his wealth, he also gets some utility from consumption.

Barberis, Huang and Santos (2001) take up this challenge. Before presenting their specification, we introduce the basic economic structure that will apply throughout our essay. At time t, the investor, whose wealth is denoted W_t, chooses a consumption level C_t and allocates his post-consumption wealth, $W_t - C_t$, across three assets. The first asset is risk-free, and earns a gross return of $R_{f,t}$ between t and $t+1$. The second asset is the stock market, which earns a gross return of $R_{S,t+1}$ over the same interval, and the third is a non-financial asset, such as human capital or housing wealth, which earns a gross return of $R_{N,t+1}$. The investor's wealth therefore evolves according to

$$W_{t+1} = (W_t - C_t)((1 - \theta_{S,t} - \theta_{N,t})R_{f,t} + \theta_{S,t}R_{S,t+1} + \theta_{N,t}R_{N,t+1}) \equiv (W_t - C_t)R_{W,t+1}, \quad (4)$$

where $\theta_{S,t}$ ($\theta_{N,t}$) is the fraction of post-consumption wealth allocated to the stock market (the non-financial asset) and $R_{W,t+1}$ is the gross return on wealth between t and $t+1$.

A stripped-down version of Barberis, Huang and Santos' (2001) framework can be written as follows. The investor maximizes

$$E_0 \sum_{t=0}^{\infty} \left[\rho^t \frac{C_t^{1-\gamma}}{1-\gamma} + b_0 \rho^{t+1} \overline{C}_t^{-\gamma} \overline{v}(G_{S,t+1}) \right], \quad (5)$$

subject to the standard budget constraint, where

$$\begin{aligned} G_{S,t+1} &= \theta_{S,t}(W_t - C_t)(R_{S,t+1} - 1) & (6) \\ \overline{v}(x) &= \begin{cases} x & \text{for } x \geq 0 \\ \lambda x & \phantom{\text{for }} x < 0 \end{cases}, \ \lambda > 1, & (7) \end{aligned}$$

and where $\overline{C}_t$ is aggregate per-capita consumption.

The first term inside the parenthesis in (5) ensures that, as in traditional models, the investor gets utility directly from consumption. The second term introduces narrow framing and loss aversion. The variable $G_{S,t+1}$ is the change in the value of the investor's stock market holdings, computed as stock market wealth at time t, $\theta_{S,t}(W_t - C_t)$, multiplied by the net stock market return, $R_{S,t+1} - 1$; $\overline{v}(G_{S,t+1})$ represents utility from this change in value. Narrow framing is therefore introduced by letting the agent get utility directly from changes in the value of just one component of his total wealth, with b_0 controlling the degree of narrow framing. Loss aversion is introduced via the piecewise linear form of $\overline{v}(\cdot)$, which makes the investor more sensitive to declines in stock market value than to increases. Finally, $\overline{C}_t^{-\gamma}$ is a neutral scaling-term that ensures stationarity in equilibrium.

Equation (6) is the simplest way of expressing the investor's exposure to stock market risk. In this case, so long as $\theta_{S,t} > 0$, a positive net return is considered a gain and, from

(7), is assigned positive utility; a negative net return is considered a loss and is assigned negative utility. Barberis, Huang and Santos (2001) work primarily with another, possibly more realistic formulation,

$$G_{S,t+1} = \theta_{S,t}(W_t - C_t)(R_{S,t+1} - R_{f,t}), \tag{8}$$

in which a stock market return is only considered a gain, and hence is only assigned positive utility, if it exceeds the risk-free rate.

In Section 2, we noted that even though narrow framing has mainly been documented in experimental settings, both the "regret" and "accessibility" interpretations suggest that people may frame the stock market narrowly as well. One could argue that they also suggest that people will frame their non-financial assets narrowly, on the grounds that the distribution of those assets' returns is also very accessible. The specification in (5) can certainly accommodate such behavior, but we have found that doing so has little effect on our results. For simplicity, then, we assume that only stock market risk is framed narrowly.

The preferences in (5) are a somewhat simplified version of Barberis, Huang and Santos' (2001) specification. In an effort to understand not only the equity premium, but also the volatility and time-series predictability of stock returns, their original model captures not only loss aversion, but also some dynamic evidence on loss aversion, sometimes known as the "house money" effect, whereby prior gains and losses affect current sensitivity to losses. The specification in (5) strips out this dynamic effect, leaving only the core features of loss aversion and narrow framing. We discuss the full model in more detail in Section 5.[7]

Barberis, Huang and Santos (2001) assign the preferences in (5), (7), and (8) to the representative agent in a simple equilibrium model, and show that when the model is calibrated to *annual* data, the narrow framing term can generate a large equity premium and a low and stable risk-free rate, even when consumption growth is smooth and only weakly correlated with stock returns. Much as in Benartzi and Thaler (1995), the intuition is that, since the investor gets direct utility from changes in the value of his stock market holdings, and is more sensitive to losses than to gains, he perceives the stock market to be very risky and will only hold the market supply if compensated by a high average return.

Of course, in assigning the utility function in (5) to a representative agent, Barberis, Huang and Santos (2001) are assuming that the key features of these preferences survive under aggregation. Intuitively, if individual investors are all loss averse over annual fluctuations in stock market wealth, it *is* hard to see why this would "wash out" in the aggregate. However, this point has not yet been formalized.

While the preference specification in (5) yields a number of insights, it also has some

[7]Barberis, Huang and Santos (2001) also consider the case in which the gains and losses are gains and losses in total wealth, rather than in stock market wealth, so that there is no narrow framing at all.

limitations. First, it does not admit an explicit value function. This makes it hard to compute attitudes to independent monetary gambles, thereby precluding us from illustrating one potential benefit of the narrow framing approach. Second, it is highly intractable in partial equilibrium settings, and so cannot be used to investigate the implications of narrow framing for portfolio choice. Finally, to ensure stationarity, the narrow framing component has to be scaled by an ad-hoc factor based on aggregate consumption.

Recently, Barberis and Huang (2004) have proposed a new preference specification that overcomes these limitations. Their starting point is recursive utility, in which the agent's time t utility, V_t, is given by

$$V_t = W(C_t, \mu(V_{t+1}|I_t)), \tag{9}$$

where $\mu(V_{t+1}|I_t)$ is the certainty equivalent of the distribution of future utility, V_{t+1}, conditional on time t information I_t, and $W(\cdot,\cdot)$ is an aggregator function that aggregates current consumption C_t with the certainty equivalent of future utility to give current utility (see Epstein and Zin, 1989, for a detailed discussion). Most implementations of recursive utility assign $W(\cdot,\cdot)$ the CES form

$$W(C,x) = ((1-\beta)C^{\rho} + \beta x^{\rho})^{\frac{1}{\rho}},\ 0 < \beta < 1,\ 0 \neq \rho < 1, \tag{10}$$

and assume homogeneity of $\mu(\cdot)$. If a certainty equivalent functional is homogeneous, it is necessarily homogeneous of degree one, so that

$$\mu(kx) = k\mu(x),\ k > 0. \tag{11}$$

In its current form, the specification in equation (9) does not allow for narrow framing: an investor with these preferences only cares about the outcome of a gamble he is offered to the extent that that outcome affects his overall wealth risk. Barberis and Huang (2004) show, however, that these preferences can be extended to accommodate narrow framing. They specify their utility function in a very general context, but for the specific three-asset setting introduced earlier, their formulation reduces to

$$V_t = W\left(C_t, \mu(V_{t+1}|I_t) + b_0 E_t(\overline{v}(G_{S,t+1}))\right), \tag{12}$$

where

$$W(C,x) = ((1-\beta)C^{\rho} + \beta x^{\rho})^{\frac{1}{\rho}},\ 0 < \beta < 1,\ 0 \neq \rho < 1 \tag{13}$$

$$\mu(kx) = k\mu(x),\ k > 0 \tag{14}$$

$$G_{S,t+1} = \theta_{S,t}(W_t - C_t)(R_{S,t+1} - R_{f,t}) \tag{15}$$

$$\overline{v}(x) = \begin{cases} x & \text{for } x \geq 0 \\ \lambda x & \text{for } x < 0 \end{cases},\ \lambda > 1. \tag{16}$$

Relative to the usual recursive specification in equation (9), this new formulation maintains the standard assumptions for $W(\cdot,\cdot)$ and $\mu(\cdot)$. The difference is that a new term, that captures loss aversion and narrow framing, has been added to the second argument of $W(\cdot,\cdot)$. As before, $G_{S,t+1}$ represents changes in the value of the investor's stock market holdings, measured relative to the risk-free rate. By letting the investor get direct utility $\overline{v}(G_{S,t+1})$ from changes in the value of this one component of his wealth, we are introducing narrow framing, with the degree of narrow framing again controlled by b_0. Loss aversion is introduced through the piecewise linearity of $\overline{v}(\cdot)$, just as in the earlier specification in (5).

Since our focus is on the effects of narrow framing, we give the certainty equivalent functional $\mu(\cdot)$ the simplest possible form, namely

$$\mu(x) = (E(x^{\zeta}))^{\frac{1}{\zeta}}, \tag{17}$$

where the exponent ζ is set to the same value as the exponent in the aggregator function, ρ. We denote this common value $1-\gamma$, so that

$$\rho = \zeta = 1-\gamma. \tag{18}$$

Throughout this article, we motivate our use of loss aversion by the fact that, as a central element of prospect theory, it is a robust feature of the way people evaluate risk in experimental settings. It is worth noting that for the preference specification in equation (12), the piecewise linear form of $\overline{v}(\cdot)$ can also be motivated without appealing to prospect theory at all, on grounds of tractability. One way to increase tractability is to impose homotheticity. Since $\mu(\cdot)$ is homogeneous of degree one, homotheticity obtains so long as $\overline{v}(\cdot)$ is *also* homogeneous of degree one. At the same time, to ensure that the first-order conditions associated with the maximization problem are both necessary and sufficient for optimality, we need $\overline{v}(\cdot)$ to be concave. The only function that is both homogeneous of degree one and concave is precisely the piecewise linear function in equation (16).

3.2 Quantitative implications

We now use the specification in equation (12) to illustrate two benefits of the narrow framing approach in more detail: first, that it can generate a high equity premium at the same time as a low and stable risk-free rate, even when consumption growth is smooth and only weakly correlated with stock returns; and then in Section 3.3, that it can do so while also making sensible predictions about attitudes to large-scale monetary gambles.

To see the first result, consider a simple economy with a representative agent who has the preferences in equation (12). As before, there are three assets. The risk-free asset is in zero net supply, and the two risky assets are each in positive net supply. Barberis and Huang (2004) show that, in this setting, the first-order conditions of optimality are

$$1 = \left[\beta R_{f,t} E_t((\frac{C_{t+1}}{C_t})^{-\gamma})\right] \left[\beta E_t((\frac{C_{t+1}}{C_t})^{-\gamma} R_{W,t+1})\right]^{\frac{\gamma}{1-\gamma}} \tag{19}$$

$$0 = \frac{E_t((\frac{C_{t+1}}{C_t})^{-\gamma}(R_{S,t+1} - R_{f,t}))}{E_t((\frac{C_{t+1}}{C_t})^{-\gamma})} + b_0 R_{f,t}(\frac{\beta}{1-\beta})^{\frac{1}{1-\gamma}}(\frac{1-\alpha_t}{\alpha_t})^{\frac{-\gamma}{1-\gamma}} E_t(\overline{v}(R_{S,t+1} - R_{f,t})) \quad (20)$$

$$0 = \frac{E_t((\frac{C_{t+1}}{C_t})^{-\gamma}(R_{W,t+1} - R_{f,t}))}{E_t((\frac{C_{t+1}}{C_t})^{-\gamma})} + b_0 R_{f,t}(\frac{\beta}{1-\beta})^{\frac{1}{1-\gamma}}(\frac{1-\alpha_t}{\alpha_t})^{\frac{-\gamma}{1-\gamma}} \theta_{S,t} E_t(\overline{v}(R_{S,t+1} - R_{f,t})) \quad (21)$$

where $\alpha_t \equiv C_t/W_t$ is the consumption-wealth ratio, and where $R_{W,t+1}$ is defined in equation (4).

We consider an equilibrium in which: (i) the risk-free rate is a constant R_f; (ii) consumption growth and stock returns are distributed as

$$\log \frac{C_{t+1}}{C_t} = g_C + \sigma_C \varepsilon_{C,t+1} \quad (22)$$

$$\log R_{S,t+1} = g_S + \sigma_S \varepsilon_{S,t+1}, \quad (23)$$

where

$$\begin{pmatrix} \varepsilon_{C,t} \\ \varepsilon_{S,t} \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho_{CS} \\ \rho_{CS} & 1 \end{pmatrix} \right), \text{ i.i.d. over time;} \quad (24)$$

(iii) the consumption-wealth ratio α_t is a constant α, which, using

$$R_{W,t+1} = \frac{W_{t+1}}{W_t - C_t} = \frac{1}{1-\alpha} \frac{C_{t+1}}{C_t}, \quad (25)$$

implies that

$$\log R_{W,t+1} = g_W + \sigma_W \varepsilon_{W,t+1}, \quad (26)$$

where

$$g_W = g_C + \log \frac{1}{1-\alpha} \quad (27)$$

$$\sigma_W = \sigma_C \quad (28)$$

$$\varepsilon_{W,t+1} = \varepsilon_{C,t+1}; \quad (29)$$

and (iv) the fraction of total wealth made up by the stock market, $\theta_{S,t}$, is a constant over time, θ_S, so that

$$\theta_{S,t} = \frac{S_t}{S_t + N_t} = \theta_S, \ \forall t, \quad (30)$$

where S_t and N_t are the total market value of the stock and of the non-financial asset, respectively. Barberis and Huang (2004) demonstrate that such a structure can indeed be embedded in a general equilibrium framework with endogeneous production.

Barberis and Huang (2004) also show that under this structure, equations (19)-(21) simplify to

$$\alpha = 1 - \beta^{\frac{1}{\gamma}} R_f^{\frac{1-\gamma}{\gamma}} e^{\frac{1}{2}(1-\gamma)\sigma_C^2} \quad (31)$$

$$
\begin{aligned}
0 &= b_0 R_f(\frac{\beta}{1-\beta})^{\frac{1}{1-\gamma}}(\frac{1-\alpha}{\alpha})^{\frac{-\gamma}{1-\gamma}}\left[e^{g_S+\frac{1}{2}\sigma_S^2}-R_f+(\lambda-1)\left[e^{g_S+\frac{1}{2}\sigma_S^2}N(\widehat{\varepsilon}_S-\sigma_S)-R_f N(\widehat{\varepsilon}_S)\right]\right]+ \\
&\quad e^{g_S+\frac{1}{2}\sigma_S^2-\gamma\sigma_S\sigma_C\rho_{CS}}-R_f \qquad (32)\\
0 &= b_0 R_f(\frac{\beta}{1-\beta})^{\frac{1}{1-\gamma}}(\frac{1-\alpha}{\alpha})^{\frac{-\gamma}{1-\gamma}}\theta_S\left[e^{g_S+\frac{1}{2}\sigma_S^2}-R_f+(\lambda-1)\left[e^{g_S+\frac{1}{2}\sigma_S^2}N(\widehat{\varepsilon}_S-\sigma_S)-R_f N(\widehat{\varepsilon}_S)\right]\right]+ \\
&\quad e^{g_C+\frac{1}{2}\sigma_C^2-\gamma\sigma_C^2}-R_f, \qquad (33)
\end{aligned}
$$

where

$$\widehat{\varepsilon}_S = \frac{\log(R_f)-g_S}{\sigma_S}. \qquad (34)$$

We use equations (31)-(33) to compute the equilibrium equity premium. First, we set the return and consumption process parameters to the values in Table 1. These values are estimated from annual data spanning the 20th century and are standard in the literature. Then, for given preference parameters β, γ, b_0 and λ, and for a given stock market fraction of total wealth θ_S, equations (31)-(33) can be solved for α, R_f and g_S, thereby giving us the equity premium.

Table 2 presents the results. We take $\beta = 0.98$ and $\theta_S = 0.2$, and consider various values of the preference parameters γ, λ and b_0. The parameter β has little effect on attitudes to risk, and our results are qualitatively very similar for a wide range of values of θ_S. The table confirms that narrow framing of stocks can generate a substantial equity premium at the same time as a low risk-free rate. For example, the parameter values $(\gamma, \lambda, b_0) = (1.5, 2, 0.1)$ produce an equity premium of 4.66% and a risk-free rate of 3.3%, while $(\gamma, \lambda, b_0) = (1.5, 3, 0.1)$ produce a premium as high as 8.16% with a risk-free rate of only 2.2%. The intuition is the same as in Benartzi and Thaler (1995) and Barberis, Huang and Santos (2001): if the agent gets utility directly from changes in the value of the stock market and, via the parameter λ, is more sensitive to losses than to gains, he perceives the stock market to be very risky and will only hold the available supply if compensated by a high average return.

Our assumption that the agent evaluates stock market gains and losses on an *annual* basis is important for our results, but not critical. Table 3 reports equity premia for an investor with the preferences in equation (12), but who evaluates stock market gains and losses at horizons other than a year. The table shows that, even though the equity premium declines as the time horizon lengthens, long evaluation periods can still generate substantial equity premia at the same time as a low risk-free rate.

The intuition for the decline in the equity premium for long evaluation periods, first pointed out by Benartzi and Thaler (1995), is straightforward. Since the distribution of stock returns has a positive mean, the probability of seeing a drop in the stock market falls as returns are aggregated at longer and longer intervals. While annual stock returns might be negative 40% of the time, five-year returns will be negative less often. An agent who is loss averse is therefore less scared of stocks when tracking their returns over longer horizons, and charges a lower equity premium.

3.3 Attitudes to large monetary gambles

We now demonstrate another attractive feature of the preference specification in equation (12), namely that it can deliver a high equity premium for parameterizations that are "reasonable," in the sense that they make sensible predictions about attitudes to independent monetary gambles. This is important because it was, in part, the difficulty researchers encountered in reconciling the equity premium with attitudes to monetary gambles that launched the equity premium literature in the first place. Economists are primarily concerned about attitudes to *large*-scale monetary gambles, so we begin with those. In Section 3.4, we also consider attitudes to small-scale gambles.

The literature has suggested a number of thought experiments involving large-scale gambles. Epstein and Zin (1990) and Kandel and Stambaugh (1991) consider an individual with wealth of $\$75,000$ and ask what premium he would pay to avoid a 50:50 chance of losing $\$25,000$ or gaining the same amount; in Kandel and Stambaugh's (1991) view, a premium of $\$24,000$ is too high, but a premium of $\$8,333$ is reasonable. Mankiw and Zeldes (1991) think about the value of x for which an agent would be indifferent between certain consumption of $\$x$ and a consumption gamble offering a 50:50 chance of $\$50,000$ consumption and $\$100,000$ consumption. Rabin (2000) suggests a mild condition, namely that an agent should accept a clearly attractive large gamble such as a 50:50 bet to win \$20 million against a $\$10,000$ loss.

It does not matter, for our results, which of these thought experiments we use. In what follows, we focus on the one suggested by Epstein and Zin (1990) and Kandel and Stambaugh (1991). In our view, a reasonable condition to impose is:[8]

Condition L: An individual with wealth of $\$75,000$ should not pay a premium higher than $\$15,000$ to avoid a 50:50 chance of losing $\$25,000$ or gaining the same amount.

Barberis and Huang (2004) show that, to avoid a gamble g offering an equal chance to win or lose x, an investor with the preferences in equation (12) would pay a premium equal to

$$\pi = \frac{A(W_t - (E(W_t + g)^{1-\gamma})^{\frac{1}{1-\gamma}}) + b_0 \frac{x}{2}(\lambda - 1)}{A + b_0 \lambda}, \tag{35}$$

where

$$A = (1-\beta)^{\frac{1}{1-\gamma}} \alpha^{-\frac{\gamma}{1-\gamma}}, \tag{36}$$

[8] We use the label "condition L" to emphasize that we are thinking about Large-scale gambles.

with α already computed in equations (31)-(33) above. In this calculation, they make the simplest possible assumption, namely that whatever degree of narrow framing b_0 and level of loss aversion λ the investor uses when thinking about stock market risk, he also uses when thinking about the independent monetary gamble g. When $b_0 = 0$, equation (35) gives the premium that would be charged by an agent with standard power utility preferences. When $b_0 > 0$, the premium in equation (35) reflects the fact that, to some extent, the investor is framing gamble g narrowly. For large b_0, equation (35) reduces to

$$\pi = \frac{x}{2\lambda}(\lambda - 1), \tag{37}$$

the premium that would be charged by an agent who evaluates gamble g completely in isolation, and who is λ times as sensitive to losses as to gains.

Using equation (35), the right-most columns in Tables 2 and 3 show, for each parameterization, the amount that the representative agent would pay, given his equilibrium holdings of risky assets, to avoid the symmetric bet in condition L. The rows in which $b_0 = 0$ reproduce a well-known result: that for power utility preferences, those values of γ low enough to make sensible predictions about attitudes to large-scale monetary gambles inevitably generate too low an equity premium.

Table 2 shows, however, that as soon as narrow framing is allowed – in other words, as soon as $b_0 > 0$ – it is easy to find parameterizations that give a high equity premium while also satisfying condition L. When $(\gamma, \lambda, b_0) = (1.5, 2, 0.1)$, for example, the investor charges a substantial equity premium of 4.66%, and a reasonable \$6,268 to avoid the $\pm$\$25,000 gamble.

How is it that the preference specification in equation (12) can reconcile attitudes to stock market risk and to the large-scale monetary gamble in condition L when other specifications have trouble doing so? To see how, note first that, in the simple representative agent economy described by conditions (i)-(iv) in Section 3.2, the equity premium is determined by the agent's attitude, in equilibrium, to adding a small amount of stock market risk to a portfolio that is only weakly correlated with the stock market. Why can we say "weakly" correlated? Since representative agent economies are calibrated to aggregate data, the correlation of stock returns and consumption growth, ρ_{CS}, must be set to a low value; when coupled with a constant consumption-wealth ratio, this immediately implies a low correlation between stock returns and returns on total wealth.[9]

To generate a substantial equity premium, then, we need the agent to be strongly averse or, at the very least, moderately averse, to a small, weakly correlated gamble. To satisfy

[9] Of course, in more general representative agent economies, the consumption-wealth ratio need not be constant, but so long as it is sufficiently stable, it should still follow that stock returns and returns on total wealth are only weakly correlated.

condition L, we need the agent to be mildly averse or, at most, moderately averse, to a large, independent gamble.

Now consider the two functions in the second argument of $W(\cdot,\cdot)$ in equation (12), namely $\mu(\cdot)$ and $\overline{v}(\cdot)$. For a γ of 1.5, the $\mu(\cdot)$ term, by virtue of its local risk-neutrality, produces only mild aversion to a small, weakly correlated gamble, but moderate aversion to a large, independent gamble. For a λ of 2, the $\overline{v}(\cdot)$ term, by virtue of being piece-wise linear, produces moderate aversion both to a small, weakly correlated gamble and to a large independent gamble. In combination, then, the two terms generate moderate aversion to a small, weakly correlated gamble – thereby giving a substantial equity premium – and moderate aversion to a large, independent gamble, thereby satisfying condition L.

3.4 Attitudes to small monetary gambles

In Section 3.3, we saw that the preferences in equation (12), capturing both loss aversion and narrow framing, can generate a large equity premium for preference parameters that also make sensible predictions about attitudes to *large*-scale monetary gambles, in that they satisfy condition L. In fact, condition L does not put very sharp restrictions on the range of equity premia that we can generate: as Table 2 shows, it can be consistent with premia as low as 0.12% or as high as 8.12%. In this section, we show that by requiring the preference specification in equation (12) to also make sensible predictions about attitudes to *small*-scale gambles, we can put much tighter bounds on the range of equity premia that narrow framing can plausibly generate.

In a way, this is not surprising. As argued earlier, in the simple representative agent economy of Section 3.2, the equity premium is determined by the agent's attitude to adding a small amount of weakly correlated stock market risk to the rest of his portfolio. If we impose constraints on the investor's attitudes to small, independent risks, it is likely that we will also constrain his attitudes to small, weakly correlated risks and thereby also, the equity premium he will charge.

What kind of condition should we impose on attitudes to small-scale gambles? As with large-scale gambles, the earlier literature has suggested a number of possible thought experiments. For consistency with our earlier discussion, we return to Epstein and Zin (1990), who ask how much an individual with wealth of $\$75,000$ would pay to avoid a 50:50 bet to lose \$250 or to win the same amount. In our view, a reasonable condition to impose here is:[10]

Condition S: An individual with wealth of $\$75,000$ should not pay a premium higher than

[10]We use the label "condition S" to emphasize that we are thinking about Small-scale gambles.

$40 to avoid a 50:50 chance of losing $250 or gaining the same amount.

Figure 1 shows how condition S sharply restricts the range of equity premia that can be generated by the preferences in equation (12). The "x" signs show, for $\gamma = 1.5$, the range of values of λ and b_0 that produce equity premia higher than 5%. Clearly, either a high sensitivity to losses λ, or a high degree of narrow framing b_0, or both, are required to generate equity premia as large as 5%. Note that our earlier condition on attitudes to large-scale gambles, condition L, is satisfied by *all* values of λ and b_0 spanned by the graph – in other words, by all pairs $(\lambda, b_0) \epsilon [0,4] \times [0,0.1]$. If condition L were the only condition constraining our choice of preference parameters, we could therefore easily obtain premia higher than 5%.

The "+" signs in the figure show the values of λ and b_0 that satisfy condition S. Imposing this condition severely restricts the range of feasible values of λ and b_0. In particular, we cannot obtain premia as high as 5% without violating it.

Even though condition S does restrict the feasible parameter set, it still allows for very sizeable equity premia. Table 4 lists some parameter values that satisfy both condition L *and* condition S, and yet still produce equity premia above 3%.

3.5 The importance of narrow framing

While narrow framing is admittedly an unusual feature of preferences, it is crucial to our results. To demonstrate this, we now show that, in the absence of narrow framing, it becomes much harder to replicate some of the attractive features of the preferences in equation (12) – much harder, for example, to reconcile a high equity premium with reasonable attitudes to large-scale monetary gambles and, in particular, with the attitudes imposed by condition L.

Consider a model in which the agent is loss averse over annual changes in *total* wealth, rather than in stock market wealth. Such a model maintains the assumptions of loss aversion and of annual evaluation of gains and losses, but by changing the focus from gains and losses in stock market wealth to gains and losses in total wealth, it removes the narrow framing. It can be written as

$$V_t = W(C_t, \mu(V_{t+1}|I_t)), \tag{38}$$

where

$$W(C, x) = ((1-\beta)C^\rho + \beta x^\rho)^{\frac{1}{\rho}}, \ 0 < \beta < 1, \ 0 \neq \rho < 1, \tag{39}$$

and where the certainty equivalent functional $\mu(\cdot)$ takes a form proposed by Gul (1991), often referred to as "disappointment aversion":

$$\mu(V)^{1-\gamma} = E(V^{1-\gamma}) + (\lambda - 1)E((V^{1-\gamma} - \mu(V)^{1-\gamma})1(V < \mu(V)), \ \gamma \neq 1. \tag{40}$$

While this specification looks somewhat messy, it is simply a function with a kink in it, that makes the investor more sensitive to losses than to gains. The parameter λ controls the relative sensitivity to losses.[11]

We consider a simple economy with a representative agent who has the preferences in equations (38)-(40). The market structure is the same as before. There are three risky assets: a risk-free asset, in zero net supply, and two risky assets, a stock market and a non-financial asset, each in positive net supply. Epstein and Zin (1989) show that the first-order conditions of optimality are:

$$\begin{aligned}
0 &= E_t\left[\phi\left(\beta^{\frac{1}{\rho}}(\frac{C_{t+1}}{C_t})^{\frac{\rho-1}{\rho}}R_{W,t+1}^{\frac{1}{\rho}}\right)\right] && (41)\\
0 &= E_t\left[\phi'\left(\beta^{\frac{1}{\rho}}(\frac{C_{t+1}}{C_t})^{\frac{\rho-1}{\rho}}R_{W,t+1}^{\frac{1}{\rho}}\right)(\frac{C_{t+1}}{R_{W,t+1}})^{\frac{\rho-1}{\rho}}(R_{S,t+1}-R_{f,t})\right] && (42)\\
0 &= E_t\left[\phi'\left(\beta^{\frac{1}{\rho}}(\frac{C_{t+1}}{C_t})^{\frac{\rho-1}{\rho}}R_{W,t+1}^{\frac{1}{\rho}}\right)(\frac{C_{t+1}}{R_{W,t+1}})^{\frac{\rho-1}{\rho}}(R_{W,t+1}-R_{f,t})\right], && (43)
\end{aligned}$$

where

$$\phi(x) = \begin{cases} \frac{x^{1-\gamma}-1}{1-\gamma} & \text{for } x \geq 1 \\ \lambda\frac{x^{1-\gamma}-1}{1-\gamma} & \phantom{\text{for }} x < 1 \end{cases}. \quad (44)$$

We look for a simple equilibrium in which conditions (i)-(iii) of Section 3.2 hold.[12] Under these conditions, equations (41)-(43) become

$$\begin{aligned}
0 &= (\frac{\beta}{1-\alpha})^{\frac{1-\gamma}{\rho}}C_1 - 1 + (\lambda-1)(\frac{\beta}{1-\alpha})^{\frac{1-\gamma}{\rho}}C_1 N(\overline{\varepsilon}_C-(1-\gamma)\sigma_C) - (\lambda-1)N(\overline{\varepsilon}_C) && (45)\\
0 &= C_2 - R_f C_3 + (\lambda-1)(C_2 N(\overline{\varepsilon}_C - \sigma_S\rho_{CS} + \gamma\sigma_C) - R_f C_3 N(\overline{\varepsilon}_C + \gamma\sigma_C)) && (46)\\
0 &= \frac{1}{1-\alpha}C_1 - R_f C_3 + (\lambda-1)(\frac{C_1}{1-\alpha}N(\overline{\varepsilon}_C-(1-\gamma)\sigma_C) - R_f C_3 N(\overline{\varepsilon}_C+\gamma\sigma_C)), && (47)
\end{aligned}$$

where

$$\begin{aligned}
C_1 &= \exp((1-\gamma)g_C + \frac{1}{2}(1-\gamma)^2\sigma_C^2) && (48)\\
C_2 &= \exp(g_S - \gamma g_C + \frac{1}{2}(\sigma_S^2 - 2\gamma\sigma_S\sigma_C\rho_{CS} + \gamma^2\sigma_C^2)) && (49)\\
C_3 &= \exp(-\gamma g_C + \frac{1}{2}\gamma^2\sigma_C^2) && (50)\\
\overline{\varepsilon}_C &= -\frac{1}{\sigma_C}(g_C + \frac{1}{\rho}\log(\frac{\beta}{1-\alpha})). && (51)
\end{aligned}$$

We use equations (45)-(47) to compute the equilibrium equity premium. As before, we set the return and consumption process parameters to the values in Table 1. Then, for given

[11] Epstein and Zin (2001) and Ang, Bekaert and Liu (2005) discuss the implementation of disappointment aversion in dynamic environments.

[12] It is straightforward to show that such a structure can be embedded in a general equilibrium model.

preference parameters β, ρ and γ, we use equation (45) to compute the consumption-wealth ratio α, equation (47) to compute the risk-free rate R_f and equation (46) to compute the mean log stock return g_S.

Since we want to check whether the parameters corresponding to any particular equity premium are reasonable – in other words, whether they satisfy condition L – we need to know the premium an agent with the preferences in equations (38)-(40) would pay to avoid a gamble to win or lose x with equal chance. Following the analysis in Epstein and Zin (1989), it can be shown that the premium π is given by

$$\frac{\pi}{W_t} = 1 - \left(\frac{(1 + \frac{x}{W_t})^{1-\gamma} + \lambda (1 - \frac{x}{W_t})^{1-\gamma}}{1 + \lambda} \right)^{\frac{1}{1-\gamma}}. \tag{52}$$

We set β and ρ, which have little effect on attitudes to risk, to 0.98 and -1, respectively. The area shaded with "+" signs in Figure 2 shows the values of γ and λ for which the representative agent satisfies condition L; in other words, the values for which, given his equilibrium holdings of risky assets and wealth of $\$75,000$, he pays a premium below $\$15,000$ to avoid a 50:50 chance of losing $\$25,000$ or winning the same amount. The area shaded with "x" signs shows the values of γ and λ for which the representative agent charges an equity premium higher than 2%. There is no overlap between the two regions: in fact, the largest equity premium that we can generate with this preference specification under condition L is 0.98%, far smaller than the equity premia derived from narrow framing in Table 2.[13]

To see the intuition for this result, recall from Section 3.3 that in the simple, representative agent economy considered here, the equity premium is determined by the agent's attitude, in equilibrium, to adding an extra dollar of stock market risk to a portfolio that is only weakly correlated with the stock market. In the absence of narrow framing, the agent evaluates this extra risk by merging it with his other risks and checking if the combination is attractive. Since the stock market is only weakly correlated with his other risks, it diversifies those other risks and so the combination *is* attractive: even a loss averse agent enjoys diversification. As a result, he charges a low equity premium. To generate a large premium, we would need to push up aversion to overall wealth risk, but this would immediately lead to a violation of condition L.

As we saw in Section 3.3., a simple way out of this difficulty is to argue that, when the agent evaluates the extra dollar of stock market risk, he does *not* fully merge it with his other risks, but rather, evaluates it in isolation; in other words, that he frames stock market risk narrowly.

[13]Epstein and Zin (1990) and Epstein and Zin (2001) obtain comparable results. See also Bekaert, Hodrick and Marshall (1997).

4 Other Applications

Barberis, Huang and Thaler (2003) argue that the preferences in equation (12) can also address a portfolio puzzle that is closely related to the equity premium puzzle, namely the stock market participation puzzle: the fact that even though stocks have a high mean return, many households have historically been unwilling to allocate any money to them. Mankiw and Zeldes (1991) report, for example, that in 1984, only 28% of households held any stock at all, and only 12% held more than $10,000 in stock. Non-participation was not simply the result of not having any liquid assets. Even among households with more than $100,000 in liquid assets, only 48% held stocks (see also Haliassos and Bertaut, 1995).

One approach to this puzzle is to argue that there are transaction costs of investing in the stock market; another is to examine whether non-stockholders have background risk that is somewhat correlated with the stock market (Heaton and Lucas 1997, 2000, Vissing-Jorgensen, 2002). A third approach is based on heterogeneity in individual preferences, and this is the one Barberis, Huang and Thaler (2003) focus on. Specifically, they show that the preferences in equation (12) can generate stock market non-participation and, mirroring our results for the equity premium, that they can do so for preference parameterizations that are reasonable, in other words, that make sensible predictions about attitudes to large-scale monetary gambles by, for example, satisfying condition L.

It is easy to see how these preferences generate non-participation: if the agent gets direct utility from fluctuations in the value of any stocks that he owns, and if he is loss averse over these fluctuations, he is naturally going to be averse to stock market risk, and may well refuse to participate.

How is it that we can generate non-participation for *reasonable* parameter values? An agent who refuses to participate in the stock market is effectively refusing to take on a small amount of a risk that is, according to Heaton and Lucas (2000), relatively uncorrelated with his other risks. To generate such attitudes at the same time as reasonable attitudes to large-scale gambles, we therefore need preferences that generate moderate aversion to a small, weakly correlated risk – thereby leading to stock market non-participation – at the same time as moderate aversion to a large independent risk, thereby satisfying condition L. As discussed in Section 3.3., the preferences in equation (12) are able to achieve exactly this.

Without narrow framing, it becomes much harder to find preference specifications that can generate non-participation for reasonable parameter values. In the absence of narrow framing, the agent decides whether to participate by mixing a small amount of stock market risk with his other risks and checking whether the combination is attractive. Since stock market risk is largely uncorrelated with his other risks, it is diversifying and so the combination is, quite generally, attractive. To prevent the agent from participating, we need to impose very high aversion to overall wealth risk, but this typically leads to implausible aversion to large-scale gambles, and in particular, to violations of condition L. This logic has been confirmed by Heaton and Lucas (2000) and Barberis, Huang and Thaler (2003), who consider a number of different specifications without narrow framing – including specifications that incorporate loss aversion – and find that all of them have trouble generating non-participation for reasonable parameter values.[14]

5 Further Extensions

5.1 Dynamic aspects of loss aversion

In a full equilibrium model, the preferences in (5) and (12) can easily deliver a high equity premium and a low and stable risk-free rate, but they have a harder time matching the empirical volatility of returns. Under these preferences, the volatility of returns is typically very similar to the volatility of dividend growth, and therefore too low.

Barberis, Huang and Santos (2001) show that incorporating *dynamic* aspects of loss aversion into the specification in (5) can help match the empirical volatility of returns.[15] Drawing on a number of different experimental tests, Thaler and Johnson (1990) argue that the degree of loss aversion is not constant over time, but depends on prior gains and losses. In particular, they present evidence that losses are less painful than usual after prior gains, perhaps because those gains cushion any subsequent loss; but that losses after prior losses are more painful than usual, perhaps because people have only limited capacity for dealing with bad news.

Barberis, Huang and Santos (2001) capture this evidence by making $\overline{v}(\cdot)$ in (5) a function not only of the current stock market return $R_{S,t+1}$ but also of *prior* gains and losses in the stock market. They then show that this raises the volatility of returns relative to the volatility of dividend growth: on good dividend news, the stock market goes up, giving the investor a cushion of prior gains and making him less sensitive to future losses; as a result, he perceives stocks to be less risky and discounts their future cash flows at a lower rate, thereby pushing prices still higher and raising the volatility of returns. The same mechanism also generates predictability in the time series: after prior gains, the investor perceives the stock market to be less risky and so pushes the price of stocks up relative to dividends; but from this point on, average returns will be lower, as the investor needs less compensation for the lower

[14] An alternative preference-based approach to the stock market participation puzzle is based on ambiguity aversion (Epstein and Schneider, 2002). This approach has some similarities to the narrow framing approach, in that it works by inducing something akin to loss aversion over the stock market gamble itself.

[15] A similar analysis can be conducted with the specification in equation (12).

perceived risk. Price-dividend ratios therefore predict returns.

One attractive feature of this mechanism is that it preserves the low correlation of stock returns and consumption growth seen in the earlier models of Section 3: since movements in the price-dividend ratio are driven by innovations to dividends, the correlation of stock returns and consumption growth is similar to the correlation of dividend growth and consumption growth, and is therefore low. This contrasts with other models of stock market volatility, such as that of Campbell and Cochrane (1999), in which movements in the price-dividend ratio are driven by innovations to consumption. These models inevitably lead to a high correlation of stock returns and consumption growth.

5.2 Other forms of narrow framing

In the economy described in Section 3, there were only two risky assets – the stock market and a non-financial asset. There were therefore only a limited number of ways in which narrow framing could manifest itself. The investor could get direct utility from stock market fluctuations, direct utility from fluctuations in the value of the non-financial asset, or both. A more realistic model would allow the investor to trade not only a broad stock market index, but individual stocks as well. Narrow framing could then, in principle, mean that the investor gets direct utility from fluctuations in the value of individual stocks that he owns. What effect would this have?

Barberis and Huang (2001) investigate this issue by extending the preferences in (5) to allow the agent to frame several assets narrowly.[16] Among other implications, they find that, if investors engage in the more extreme form of narrow framing whereby they frame even individual stocks narrowly, the equity premium can be even higher than in the case studied in Section 3, where they frame only their overall *portfolio* of stocks narrowly: if investors worry about fluctuations in highly volatile individual stocks rather than just in the less volatile aggregate stock market, they perceive stocks to be very risky and charge a very high premium for holding stock in equilibrium.

Is it plausible that people might frame individual stocks narrowly? From a theoretical perspective, it is hard to tell. Consider Kahneman's (2003) "accessibility" theory of framing. It is true that for most investors, information about the value of individual stocks that they own is highly accessible. But so too is information about the value of their overall stock portfolio, and it seems that given a choice between the broader frame and the narrower one, people will choose the normatively more acceptable frame, namely the broader one, for their decision-making.

[16] A similar analysis can be performed using the formulation in equation (12).

Under the alternative theory that narrow framing is related to non-consumption utility such as regret, framing at the level of individual stocks becomes more plausible. If one of the investor's stocks performs poorly, he may regret the specific decision to buy that stock. Gains and losses on individual stocks can therefore be carriers of utility in their own right, and the investor may take this into account when making decisions.

The framing of individual stocks is also supported by the well-known disposition effect – the fact that when individual investors sell stocks in their portfolios, they tend to sell stocks that have gone up in value since they bought them, rather than stocks that have gone down (Shefrin and Statman, 1984, Odean, 1998). The leading explanation of this finding is that people get direct utility from realizing a loss on an individual stock that they own and that this leads them to postpone selling a losing stock for as long as possible.[17]

6 Conclusion and Future Directions

In this essay, we discuss a recent approach to the equity premium puzzle. The broad theme of this approach is that we may be able to improve our understanding of how people evaluate stock market risk by looking at how they evaluate risk in experimental settings. More specifically, this approach argues that loss aversion and narrow framing, two of the most important ideas to emerge from the experimental literature on decision-making under risk, may also play an important role in the stock market setting.

We discuss various ways of incorporating loss aversion and narrow framing into more traditional utility functions, and then show that models with these features may indeed offer an attractive way of thinking about the equity premium puzzle. For example, they can generate a high equity premium and a low and stable risk-free rate, even when consumption growth is smooth and only weakly correlated with the stock market; moreover, they can do this for parameter values that are reasonable, in that they make sensible predictions about attitudes to independent monetary gambles. A parallel result holds in the case of the stock market participation puzzle, a portfolio puzzle that is closely related to the equity premium puzzle: with narrow framing, we can generate non-participation for very reasonable parameter values.

While these initial results are promising, much work remains to be done. The most obvious direction for future research is to think about other testable implications of the

[17]Tax considerations point to the selling of prior losers, so they cannot explain the disposition effect. Nor can the effect be explained by a rational belief in mean-reversion: the stocks that individual investors sell actually outperform the ones they buy (Odean, 1998). Even an irrational belief in mean-reversion is an unlikely explanation because individual investors *not* holding stocks that have recently gone down in value do not tend to buy them, as they should if they were anticipating a rebound.

loss aversion / narrow framing view. For example, while narrow framing makes a blanket prediction of non-participation in the stock market, does it also make more detailed predictions about what kinds of people are more likely to participate than others? Does it predict changes in participation over time, perhaps due to changes in framing? Are there any real-world situations in which people are asked to make a certain financial decision after seeing some data, and that have the feature that while everyone sees the same data, some people see it presented somewhat differently than others? The differences in the way the data is presented could lead people to frame future outcomes differently, and therefore to make different choices.

Researchers have already begun testing our view of the equity premium and participation puzzles. Dimmock (2005) describes a recent survey in Holland in which subjects were given a decision problem involving riskless choice. Responses to this decision problem can be used to extract estimates of individual loss aversion. After extracting these estimates from the data, Dimmock (2005) finds that individuals with greater loss aversion are indeed less likely to participate in the stock market.

Narrow framing is harder to measure than loss aversion, but successful tests of narrow framing in other settings suggest that progress can also be made in the context of the equity premium and stock market participation puzzles. Kumar and Lim (2004), for example, test the idea that narrow framing is behind the disposition effect by checking whether individual investors who engage less in narrow framing also exhibit less of a disposition effect. They identify these investors as those who tend to execute more than one trade on any given day, and who therefore might pay less attention to the outcome of any one transaction. They find that these investors do indeed exhibit less of a disposition effect.

Our attempt to bring psychology into economics has also served to highlight some areas of the original psychology where more research would be valuable. While there is now ample evidence that, in some situations, people frame narrowly, we still do not fully understand when people frame narrowly and when they do not, nor what the underlying causes of narrow framing are. Similarly, while loss aversion itself is a robust and well-documented phenomenon, much less is known about its dynamic aspects: for example, about how past gains and losses affect subsequent loss aversion. Thaler and Johnson (1990) provide some valuable evidence on this point, but it is hard to believe that theirs is the last word. A better understanding of these issues, perhaps through more experimental research, may eventually help us craft better models of how people evaluate stock market risk.

7 References

Ang, A., Bekaert G., and J. Liu (2005), "Why Stocks May Disappoint," *Journal of Financial Economics* 76, 471-508.

Barberis, N. and M. Huang (2001), "Mental Accounting, Loss Aversion and Individual Stock Returns," *Journal of Finance* 56, 1247-1292.

Barberis, N. and M. Huang (2004), "Preferences with Frames: A New Utility Specification that Allows for the Framing of Risks," Working paper, Yale University.

Barberis, N., Huang M., and T. Santos (2001), "Prospect Theory and Asset Prices," *Quarterly Journal of Economics* 116, 1-53.

Barberis, N., Huang M., and R. Thaler (2003), "Individual Preferences, Monetary Gambles and Stock Market Participation," Working paper, Yale University.

Bekaert, G., Hodrick R., and D. Marshall (1997), "The Implications of First-order Risk Aversion for Asset Market Risk Premiums," *Journal of Monetary Economics* 40, 3-39.

Benartzi, S. and R. Thaler (1995), "Myopic Loss Aversion and the Equity Premium Puzzle," *Quarterly Journal of Economics* 110, 75-92.

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

Bossaerts, P., Plott, C., and W. Zame (2003), "Prices and Portfolio Choices in Financial Markets: Theory and Experiment," Working paper, California Institute of Technology.

Campbell, J.Y. and J. Cochrane (1999), "By Force of Habit: A Consumption-based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205-251.

Cicchetti C. and J. Dubin (1994), "A Microeconometric Analysis of Risk Aversion and the Decision to Self Insure," *Journal of Political Economy* 102, 169-186.

Dimmock, S. (2005), "Loss Aversion and Household Portfolio Choice," Working paper, Michigan State University.

Epstein, L. and M. Schneider (2002), "Learning Under Ambiguity," Working paper, University of Rochester.

Epstein, L. and S. Zin (1989), "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Perspective," *Econometrica* 57, 937-969.

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

Epstein, L. and S. Zin (2001), “The Independence Axiom and Asset Returns,” *Journal of Empirical Finance* 8, 537-572.

Gneezy, U. and J. Potters (1997), “An Experiment on Risk Taking and Evaluation Periods,” *Quarterly Journal of Economics* 112, 631-645.

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

Grgeta A. and R. Thaler (2003), “Estimating Risk Aversion from the Choice of Deductibles in Automobile Insurance,” Working paper, University of Chicago.

Gul, F. (1991), “A Theory of Disappointment Aversion,” *Econometrica* 59, 667-686.

Haliassos, M. and C. Bertaut (1995), “Why Do So Few Hold Stocks?” *Economic Journal* 105, 1110-1129.

Heaton, J. and D. Lucas (1997), “Market Frictions, Savings Behavior and Portfolio Choice,” *Macroeconomic Dynamics* 1, 76-101.

Heaton, J. and D. Lucas (2000), “Portfolio Choice in the Presence of Background Risk,” *Economic Journal* 110, 1-26.

Kahneman, D. (2003), “Maps of Bounded Rationality: Psychology for Behavioral Economics,” *American Economic Review* 93, 1449-1475.

Kahneman, D., Knetsch A. and R. Thaler (1990), “Experimental Tests of the Endowment Effect and the Coase Theorem,” *Journal of Political Economy* 98, 1325-1348.

Kahneman, D. and D. Lovallo (1993), “Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking,” *Management Science* 39, 17-31.

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

Kahneman, D. and A. Tversky (1983), “Choices, Values and Frames,” *American Psychologist* 39, 341-350.

Kandel, S. and R. Stambaugh (1991), “Asset Returns and Intertemporal Preferences,” *Journal of Monetary Economics* 27, 39-71.

Kumar, A. and S. Lim (2004), "One Trade at a Time: Narrow Framing and Stock Investment Decisions of Individual Investors," Working paper, University of Notre Dame.

Mankiw, N.G. and S. Zeldes (1991), "The Consumption of Stockholders and Non-stockholders," *Journal of Financial Economics* 29, 97-112.

Mehra, R. and E. Prescott (1985), "The Equity Premium: A Puzzle," *Journal of Monetary Economics* 15, 145-161.

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.

Rabin, M. and R. Thaler (2001), "Anomalies: Risk Aversion," *Journal of Economic Perspectives* 15, 219-232.

Read, D., Loewenstein G., and M. Rabin (1999), "Choice Bracketing," *Journal of Risk and Uncertainty* 19, 171-197.

Redelmeier, D. and A. Tversky (1992), "On the Framing of Multiple Prospects," *Psychological Science* 3, 191-193.

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

Thaler, R. (1980), "Towards a Positive Theory of Consumer Choice," *Journal of Economic Behavior and Organization* 1, 39-60.

Thaler, R. and E. Johnson (1990), "Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice," *Management Science* 36, 643-660.

Thaler, R., Tversky A., Kahneman D., and A. Schwartz (1997), "The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test," *Quarterly Journal of Economics* 112, 647-661.

Tversky, A. and D. Kahneman (1981), "The Framing of Decisions and the Psychology of Choice," *Science* 211, 453-458.

Tversky, A. and D. Kahneman (1986), "Rational Choice and the Framing of Decisions," *Journal of Business* 59, 5251-78.

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

Vissing-Jorgensen, A. (2002), "Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures," Working paper, Northwestern University.

Table 1: Parameter values for a representative agent equilibrium model: g_C and σ_C are the mean and standard deviation of log consumption growth, σ_S is the standard deviation of log stock returns, and ρ_{CS} is the correlation of log consumption growth and log stock returns.

Parameter	
g_C	1.84%
σ_C	3.79%
σ_S	20.0%
ρ_{CS}	0.10

Table 2: The table shows, for given aversion to consumption risk γ, sensitivity to narrowly framed losses λ, and degree of narrow framing b_0, the risk-free rate R_f and equity premium EP generated by narrow framing in a simple representative agent economy. π_L is the premium the representative agent would pay, given his equilibrium holdings of risky assets and wealth of $75,000, to avoid a 50:50 bet to win or lose $25,000.

γ	λ	b_0	R_f	EP	π_L
1.5	2	0	4.7%	0.12%	$6,371
1.5	2	0.05	3.6%	3.75%	$6,285
1.5	2	0.1	3.3%	4.66%	$6,268
1.5	3	0	4.7%	0.12%	$6,371
1.5	3	0.05	2.6%	7.06%	$8,037
1.5	3	0.1	2.2%	8.16%	$8,193
3	2	0	6.9%	0.24%	$11,754
3	2	0.05	4.9%	3.40%	$8,214
3	2	0.1	4.3%	4.45%	$7,305
3	3	0	6.8%	0.24%	$11,754
3	3	0.05	2.8%	6.91%	$8,901
3	3	0.1	2.0%	8.22%	$8,561

Table 3: The table shows, for given aversion to consumption risk γ, sensitivity to narrowly framed losses λ, and degree of narrow framing b_0, the risk-free rate R_f and equity premium EP generated by narrow framing in a simple representative agent economy. π_L is the premium the representative agent would pay, given his equilibrium holdings of risky assets and wealth of $75,000, to avoid a 50:50 bet to win or lose $25,000. T is the horizon, in years, over which stock market gains and losses are measured.

T	γ	λ	b_0	R_f	EP	π_L
0.5	1.5	2	0.1	2.4%	7.60%	$6,257
1	1.5	2	0.1	3.3%	4.66%	$6,268
2	1.5	2	0.1	3.9%	2.56%	$6,287
3	1.5	2	0.1	4.1%	1.72%	$6,300

Table 4: The table shows, for given aversion to consumption risk γ, sensitivity to narrowly framed losses λ, and degree of narrow framing b_0, the risk-free rate R_f and equity premium EP generated by narrow framing in a simple representative agent economy. π_L (π_S) is the premium the representative agent would pay, given his equilibrium holdings of risky assets and wealth of $75,000, to avoid a 50:50 bet to win or lose $25,000 ($250).

γ	λ	b_0	R_f	EP	π_L	π_S
1.5	2	0.035	3.7%	3.18%	$6,296	$39.1
1.5	3	0.012	3.8%	3.12%	$7,268	$38.4
3	2	0.042	5.1%	3.10%	$8,494	$37.5
3	3	0.016	5.0%	3.26%	$10,184	$38.9

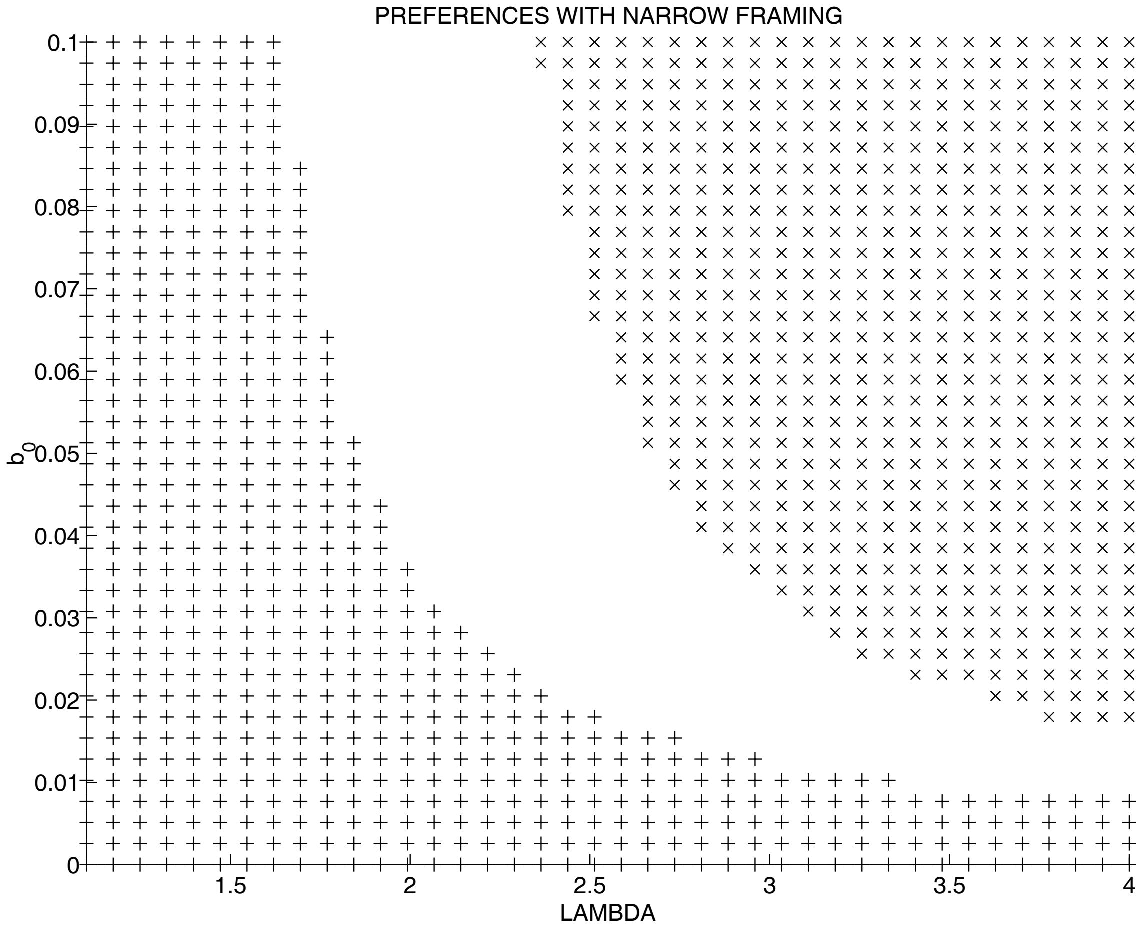

Figure 1. The "x" signs show the parameter values for which an agent who is loss averse over stock market risk would charge an equity premium higher than 2% in a simple representative agent economy. The "+" signs show where the agent would pay a premium below $40 to avoid a 50:50 bet to win or lose $250 at a wealth level of $75,000.

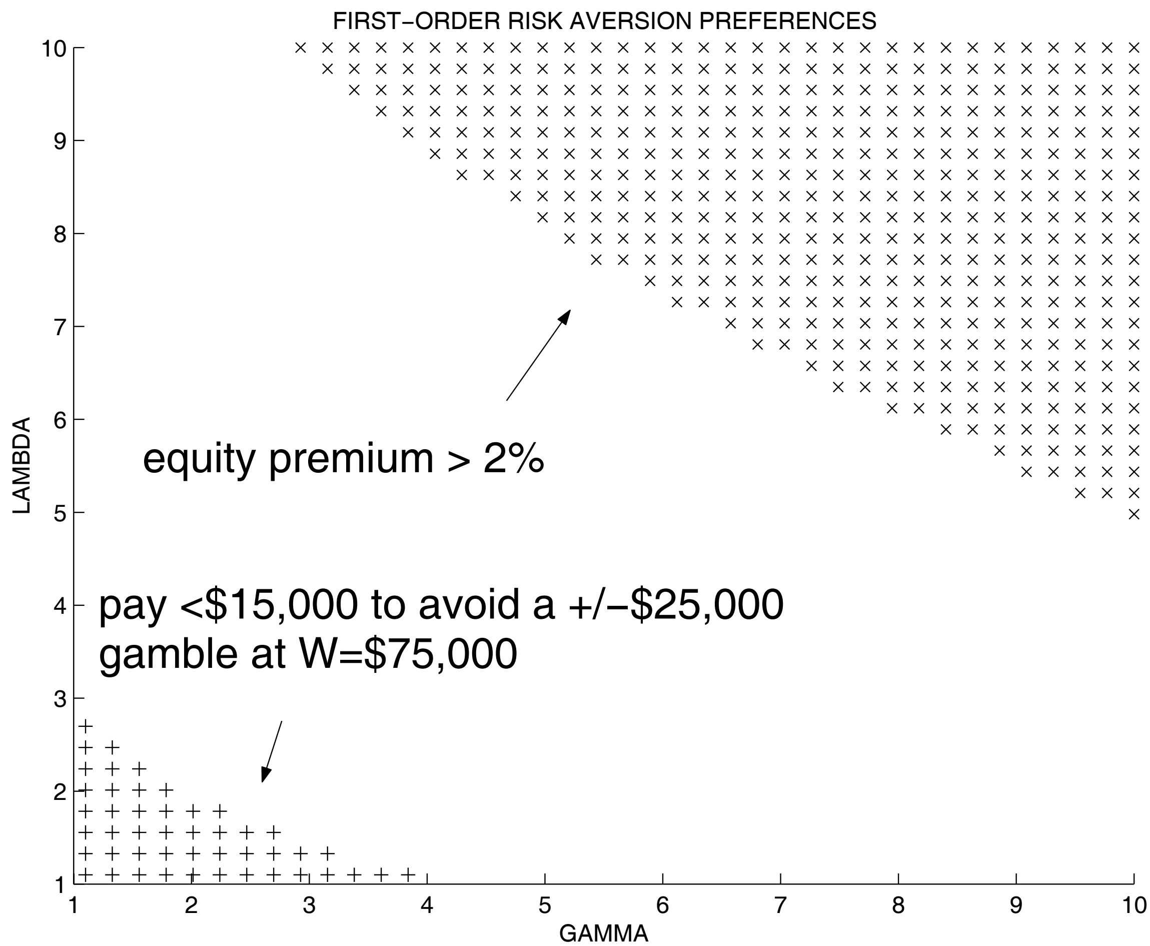

Figure 2. The "x" signs show the parameter values for which an agent with a recursive utility function with Gul (1991)-type certainty equivalent would charge an equity premium higher than 2% in a simple representative agent economy. The "+" signs show where the agent would pay a premium below $15,000 to avoid a 50:50 bet to win or lose $25,000 at a wealth level of $75,000.