

Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing

Nicholas Barberis, Ming Huang, and Richard H. Thaler
Yale University, Cornell University, and University of Chicago

April 2006*

Abstract

We argue that “narrow framing,” whereby an agent who is offered a new gamble evaluates that gamble in isolation, separately from other risks she already faces, may be a more important feature of decision-making than previously realized. Our starting point is the evidence that people are often averse to a small, independent gamble, even when the gamble is actuarially favorable. We find that a surprisingly wide range of utility functions, including many non-expected utility specifications, have trouble explaining this evidence; but that this difficulty can be overcome by allowing for narrow framing. Our analysis makes predictions as to what kinds of preferences can most easily address the stock market participation puzzle, as well as other related financial puzzles. We confirm these predictions in a simple portfolio choice setting.

JEL classification: D1, D8, G11, G12

Keywords: risk aversion, framing, loss aversion, stock market participation

*We thank Douglas Bernheim and Vincent Crawford (the editors), three anonymous referees, John Campbell, Darrell Duffie, Larry Epstein, John Heaton, Daniel Kahneman, David Laibson, Erzo Luttmer, Terrence Odean, Mark Rubinstein, Bill Sharpe, Jeffrey Zwiebel, and seminar participants at Duke University, Emory University, Harvard University, INSEAD, London Business School, London School of Economics, Massachusetts Institute of Technology, New York University, Northwestern University, Princeton University, Stanford University, University of California at Berkeley, University of Chicago, University of Pennsylvania, Yale University, and the NBER for helpful comments.

1 Introduction

Economists, and financial economists in particular, have long been interested in how people evaluate risk. In this paper, we try to shed new light on this topic. Specifically, we argue that a feature known as “narrow framing” may play a more important role in decision-making under risk than previously realized. In traditional models, which define utility over total wealth or consumption, an agent who is offered a new gamble evaluates that gamble by merging it with the other risks she already faces and checking whether the combination is attractive. Narrow framing, by contrast, occurs when an agent who is offered a new gamble evaluates that gamble to some extent in isolation, separately from her other risks.

Our starting point is the evidence that people are often reluctant to take a small, independent gamble, even when the gamble is actuarially favorable. For example, in an experimental setting, Tversky and Kahneman (1992) find that approximately half of their subjects turn down a small gamble with two equiprobable outcomes, even though the potential gain is twice the potential loss.

In this paper, we try to understand what kinds of preferences can explain such attitudes to risk. It is already understood that, without a property known as “first-order” risk aversion, it is hard for preferences to explain the rejection of a small, independent, actuarially favorable gamble (Epstein and Zin 1990, Rabin 2000). Our own contributions are: first, to show that, surprisingly, even utility functions that do exhibit first-order risk aversion have difficulty explaining why such a gamble would be rejected; and, second, that utility functions that combine first-order risk aversion with narrow framing have a much easier time doing so.¹

The intuition for these results is straightforward. Suppose that an agent with first-order risk aversion, but who does *not* engage in narrow framing, is offered a small, independent, actuarially favorable gamble. Suppose also, as is reasonable, that the agent faces some pre-existing risk, such as labor income risk or house price risk. In the absence of narrow framing, the agent must evaluate the new gamble by merging it with her pre-existing risk and checking if the combination is attractive. It turns out that the combination *is* attractive: since the new gamble is independent of the agent’s other risks, it diversifies those other risks, and, even though first-order risk averse, the agent finds this useful. The only way to make the agent reject the gamble is to choose an extreme parameterization of her utility function. However, such a parameterization typically implies, counterfactually, the rejection of highly attractive gambles with larger stakes.

A simple way out of this difficulty, then, is to suppose that, when the agent evaluates

¹We give the formal definition of first-order risk aversion in Section 2. In informal terms, a utility function exhibits first-order risk aversion if it is locally risk averse, unlike most standard preferences, which are locally risk-neutral; loss aversion, whereby the utility function is kinked at the current wealth level, is an example of first-order risk aversion.

the small gamble, she does *not* fully merge it with her pre-existing risk, but rather, thinks about it in isolation, to some extent; in other words, she frames the gamble narrowly. Using a recently developed preference specification that allows for both first-order risk aversion *and* narrow framing, we confirm that a utility function with these features can easily explain the rejection of a small, independent, actuarially favorable gamble, while also making sensible predictions about attitudes to large-scale gambles: if the agent’s first-order risk aversion is focussed directly on the small gamble rather than just on overall wealth risk, she will be reluctant to take it.

Our analysis of independent monetary gambles has useful implications for financial markets; for example, for understanding what kinds of preferences can address the stock market participation puzzle: the fact that, even though the stock market has a high mean return and a low correlation with other household risks, many households have historically been reluctant to allocate any money to it (Mankiw and Zeldes 1991, Haliassos and Bertaut 1995).

It is already understood that, without first-order risk aversion, it is hard to find a preference-based explanation of this evidence. Our analysis of small, independent gambles suggests a more surprising prediction: that even preferences that do exhibit first-order risk aversion will have trouble explaining non-participation. It also suggests that preferences that combine first-order risk aversion with narrow framing – in this case, narrow framing of stock market risk – will have an easier time doing so. We confirm these predictions in a simple portfolio choice setting.

While the term “narrow framing” was first used by Kahneman and Lovallo (1993), the more general concept of “decision framing” was introduced much earlier by Tversky and Kahneman (1981). There are already several cleverly-designed laboratory demonstrations of narrow framing in the literature.² This paper shows that a more basic piece of evidence on attitudes to risk, not normally associated with narrow framing, may also need to be thought of in these terms. Moreover, while existing examples of narrow framing do not always have obvious counterparts in the everyday risks people face, the simple risks we consider *do* – not least in stock market risk – making our results applicable in a variety of contexts.

The idea that a combination of first-order risk aversion and narrow framing may be relevant for understanding aversion to small gambles has been proposed before, most notably by Rabin and Thaler (2001). Those authors do not, however, demonstrate the sense in which first-order risk aversion, on its own, is insufficient. This is the crucial issue we tackle here.

If narrow framing does, sometimes, play a role in the way people evaluate risk, it is important that we learn more about its underlying causes. At the end of the paper, we discuss some interpretations of narrow framing, including the possibility that it arises when

²Read, Loewenstein, and Rabin (1999) survey many examples of narrow framing, including those documented by Tversky and Kahneman (1981) and Redelmeier and Tversky (1992).

people make decisions intuitively, rather than through effortful reasoning (Kahneman 2003). We argue that the recent attempts to outline a theory of narrow framing have made framing-based hypotheses much more testable than they were before.

2 Attitudes to Monetary Gambles

People are often averse to a small, independent gamble, even when the gamble is actuarially favorable. In an experimental setting, Tversky and Kahneman (1992) find that approximately half of their subjects turn down a small gamble with two equiprobable outcomes, even though the potential gain is twice the potential loss. In a more recent study, Abdellaoui, Bleichrodt, and Paraschiv (2005) also find that their median subject is indifferent to a small gamble with two equiprobable outcomes only when the gain is twice the potential loss.³

In this paper, we try to understand what kinds of preferences can explain an aversion to a small, independent, actuarially favorable gamble. For concreteness, we consider a small, independent gamble whose potential gain is much less than twice the potential loss, namely⁴

$$G_S = (550, \frac{1}{2}; -500, \frac{1}{2}),$$

to be read as “gain \$550 with probability $\frac{1}{2}$ and lose \$500 with probability $\frac{1}{2}$,” and ask: What kinds of preferences can explain the rejection of this gamble? We do not insist that a utility function be able to explain the rejection of G_S at all wealth levels, but rather that it do so over a range of wealth levels – at wealth levels below \$1,000,000, say. We take “wealth” to mean total wealth, including both financial assets and non-financial assets such as human capital.

Is it reasonable to posit that individuals tend to reject G_S even at a wealth level of \$1,000,000, a wealth level that probably exceeds that of Tversky and Kahneman’s (1992) median subject? In Barberis, Huang, and Thaler (2003), an earlier version of the current paper, we ask four groups of people – 68 part-time MBA students, 30 financial advisors, 19 Chief Investment Officers at large asset management firms, and 34 clients of a U.S.

³This evidence is not necessarily inconsistent with risk-seeking behavior like casino gambling or the buying of lottery tickets. Lottery tickets are different from Tversky and Kahneman’s (1992) gambles, in that they offer a tiny probability of substantial gain, rather than two equiprobable outcomes. An individual can be averse to a 50:50 bet offering a gain against a loss, even if she is risk-seeking over low probability gains. Gambling is a special phenomenon, in that people would not accept the terms of trade offered at a casino if those terms were offered by their bank, say. It must be that, in the casino setting, people either misestimate their chance of winning or else receive utility from the gambling activity itself.

⁴The “S” subscript in G_S stands for Small stakes. We sometimes use the notation X/Y to refer to a 50:50 bet to win \$ X or lose \$ Y . G_S is therefore a “550/500” bet.

bank’s private wealth management division (the median wealth in this last group exceeds \$10 million) – about their attitudes to a hypothetical G_S . In all four groups, the majority reject G_S , and even in the wealthiest group, the rejection rate is 71%. Playing G_S for real money with a second group of MBA students leads to an even higher rejection rate than among the first group of MBAs.

We focus on *preference*-based explanations for aversion to small gambles because the alternative explanations seem incomplete. One alternative view is that the rejection of gambles like G_S is due to transaction costs that might be incurred if a liquidity-constrained agent has to finance a loss by selling illiquid assets (Chetty 2003). Such a mechanism does not, however, explain why wealthy subjects with substantial liquid assets are often averse to G_S . Nor is suspicion a satisfactory explanation for the rejection of G_S : the fear, for example, that the experimenter is using a biased coin to determine the outcome. Offering to use a subject’s coin instead does little to change attitudes to the gamble.⁵

Of course, many utility functions can capture an aversion to G_S by simply assuming high risk aversion. To ensure that the assumed risk aversion is realistic, we insist that the utility functions we consider also make sensible predictions about attitudes to *large*-scale gambles by, for example, predicting acceptance of the large, independent gamble⁶

$$G_L = (20,000,000, \frac{1}{2}; -10,000, \frac{1}{2})$$

over a reasonable range of wealth levels – wealth levels above \$100,000, say.

In summary, then, we are interested in understanding what kinds of preference specifications can predict both:

- I. that G_S is rejected for wealth levels $W \leq \$1,000,000$, and
- II. that G_L is accepted for wealth levels $W \geq \$100,000$.

For certain utility specifications, it can make a difference whether gambles are “immediate” or “delayed.” A gamble is immediate if its uncertainty is resolved at once, before any further consumption decisions are made. A delayed gamble, on the other hand, might play out as follows: in the case of G_S , the subject is told that, at some point in the next few months, she will be contacted and informed either that she has just won \$550 or that she has lost \$500, the two outcomes being equally probable and independent of other risks.

⁵It is also unlikely that people turn down G_S simply because they do not want, or feel able, to evaluate it. In our experience, subjects, when debriefed, typically explain their aversion to G_S by saying that “the pleasure of a \$550 gain doesn’t compensate for the pain of a \$500 loss.” It therefore appears that they *do* evaluate the gamble, but find it unattractive.

⁶The “ L ” subscript in G_L stands for Large stakes.

Although certain utility functions can predict different attitudes to immediate and delayed gambles, people do not appear to treat the two kinds of bets very differently. Barberis, Huang, and Thaler (2003), for example, record very similar rejection rates for immediate and delayed versions of G_S . We therefore look for preference specifications that can satisfy conditions I and II both when G_S and G_L are immediate, *and* when they are delayed. Since a delayed gamble’s uncertainty is resolved only after today’s consumption is set, it must be analyzed in a multiperiod framework. We therefore work with intertemporal preferences, not static ones, throughout the paper.

2.1 Utility Functions

To structure our discussion, we introduce a simple taxonomy of intertemporal preferences. Our list is not exhaustive, but it covers most of the preference specifications used in economics. In parentheses, we list the abbreviations that we use to refer to specific classes of utility functions.

[EU]: Expected utility preferences

Non-expected utility preferences:

[R-EU]: Recursive utility with EU certainty equivalent

[R-SORA]: Recursive utility with non-EU, second-order risk averse certainty equivalent

[R-FORA]: Recursive utility with non-EU, first-order risk averse certainty equivalent

Expected utility preferences are familiar enough. In an intertemporal setting, non-expected utility is typically implemented using a recursive structure in which time t utility, V_t , is defined through

$$V_t = W(C_t, \mu(\tilde{V}_{t+1}|I_t)), \quad (1)$$

where $\mu(\tilde{V}_{t+1}|I_t)$ is the certainty equivalent of the distribution of future utility \tilde{V}_{t+1} conditional on time t information I_t , and $W(\cdot, \cdot)$ is an aggregator function which aggregates current consumption C_t with the certainty equivalent of future utility to give current utility.

The three kinds of recursive utility on our list differ in the properties they impose on the certainty equivalent functional $\mu(\cdot)$. In the first kind, labelled “R-EU,” $\mu(\cdot)$ has the expected utility form, so that

$$\mu(\tilde{X}) = h^{-1} E h(\tilde{X}), \quad (2)$$

for some increasing $h(\cdot)$. In the “R-SORA” class, $\mu(\cdot)$ is non-EU but exhibits “second-order” risk aversion, which means that it predicts risk-neutrality for infinitesimal risks. In simple terms, utility functions with second-order risk aversion are smooth. Finally, we consider the

“R-FORA” class, in which $\mu(\cdot)$ is non-EU and exhibits “first-order” risk aversion, which, as noted earlier, means that it predicts risk aversion even over infinitesimal bets.⁷

3 The Limitations of Standard Preferences

It is already understood that the first three kinds of utility functions on our list – the EU, R-EU, and R-SORA classes – have trouble satisfying condition I (the rejection of G_S) while also satisfying condition II (the acceptance of G_L). Rabin (2000), for example, shows that no increasing, concave utility function in the EU class is consistent with both conditions. The intuition is straightforward. An individual with EU preferences is locally risk-neutral; since G_S involves small stakes, she would normally take it. To get her to reject it, in accordance with condition I, we need very high local curvature. In fact, her utility function must have high local curvature at *all* wealth levels below \$1,000,000, because condition I requires rejection of G_S at all points in that range. Rabin’s (2000) analysis shows that “linking together” all of these locally concave pieces gives a utility function with global risk aversion so high that the agent rejects even the very favorable large gamble G_L .

While R-EU preferences are non-expected utility, the fact that, in this case, the certainty equivalent functional $\mu(\cdot)$ is in the expected utility class means that an agent with these preferences evaluates risk using an expected utility function. Rabin’s (2000) argument therefore also applies to R-EU preferences: no utility function in this class is consistent with both conditions I and II.

R-SORA preferences can, in principle, satisfy conditions I and II, but only with difficulty. An agent with R-SORA preferences is locally risk-neutral and will therefore normally be happy to accept a small, actuarially favorable gamble like G_S . To make her reject it, we need very high local curvature, which, in turn, requires an extreme parameterization. However, such a parameterization almost always implies high global risk aversion, thereby making the individual reject attractive, large-scale gambles like G_L .⁸

⁷The labels “second-order” and “first-order” are due to Segal and Spivak (1990). Formally, second-order risk aversion means that the premium paid to avoid an actuarially fair gamble $k\tilde{\varepsilon}$ is, as $k \rightarrow 0$, proportional to k^2 . Under first-order risk aversion, the premium is proportional to k . While non-EU functions can exhibit either first- or second-order risk aversion, utility functions in the EU class can generically only exhibit second-order risk aversion: an increasing, concave utility function can only have a kink at a countable number of points.

⁸Those R-SORA preferences that are able to satisfy conditions I and II typically stumble on the following additional observation: that people tend to reject $1.1y/y$ and accept $4y/y$ for a wide range of values of y . R-SORA preferences find it hard to explain such “linear” attitudes, as they need to invoke very strong *non*-linearity, or local curvature, to capture the rejection of $1.1y/y$ for just one value of y . For more discussion of EU, R-EU, and R-SORA preferences, see Barberis, Huang, and Thaler (2003).

The main contribution of this paper is to show that even the fourth preference class on our list – the R-FORA class, in which the certainty equivalent functional $\mu(\cdot)$ is non-EU and first-order risk averse – has difficulty satisfying conditions I and II. This is a surprising result because, at first sight, it appears that R-FORA preferences can easily be consistent with these conditions. In fact, as noted by Ang, Bekaert, and Liu (2005) and anticipated even earlier by Epstein and Zin (1990), such preferences have no trouble satisfying conditions I and II, *so long as the gambles are played out immediately*, a caveat that will turn out to be critical.

To see what happens when the gambles are immediate, consider the following R-FORA preferences:

$$W(C, \mu) = ((1 - \beta)C^\rho + \beta\mu^\rho)^{\frac{1}{\rho}}, \quad 0 \neq \rho < 1, \quad (3)$$

where $\mu(\cdot)$ takes a form developed by Gul (1991),

$$\mu(\tilde{V})^{1-\gamma} = E(\tilde{V}^{1-\gamma}) + (\lambda - 1)E((\tilde{V}^{1-\gamma} - \mu(\tilde{V})^{1-\gamma})1(\tilde{V} < \mu(\tilde{V}))), \quad 0 < \gamma \neq 1, \lambda > 1. \quad (4)$$

These preferences are often referred to as “disappointment aversion” preferences: the agent incurs disutility if the outcome of \tilde{V} falls below the certainty equivalent $\mu(\tilde{V})$. The parameter λ governs the degree of disutility, in other words, how sensitive the agent is to losses as opposed to gains. A $\lambda > 1$ puts a kink in the utility function at the certainty equivalent point, thereby generating first-order risk aversion.

Epstein and Zin (1989), who give a careful exposition of recursive utility, propose that, to evaluate an immediate gamble \tilde{v} at time τ , an agent with the recursive preferences in (1) inserts an infinitesimal time step $\Delta\tau$ immediately before time τ consumption C_τ is chosen, and then applies the recursive calculation over this time step, checking whether the utility from taking the gamble,

$$W(0, \mu(\tilde{V}_{\tau+\Delta\tau})) = W(0, \mu(J(\tilde{W}_{\tau+\Delta\tau}))) = W(0, \mu(J(W_\tau + \tilde{v}))), \quad (5)$$

where $J(\cdot)$ is the agent’s value function at time $\tau + \Delta\tau$, is greater than the utility from *not* taking the gamble,

$$W(0, \mu(V_{\tau+\Delta\tau})) = W(0, \mu(J(W_{\tau+\Delta\tau}))) = W(0, \mu(J(W_\tau))). \quad (6)$$

The decision therefore comes down to comparing $\mu(J(W_\tau + \tilde{v}))$ and $\mu(J(W_\tau))$.

Suppose that, aside from \tilde{v} , the agent’s other investment opportunities are i.i.d. across periods, and that the outcome of \tilde{v} is independent of, and does not affect, these other opportunities. It is then straightforward to show that, for $t \geq \tau + \Delta\tau$, the time t value function in the case of (3)-(4) takes the form

$$J(W_t) = AW_t \quad (7)$$

for some constant A . Equations (5) and (6) immediately imply that the agent evaluates \tilde{v} by comparing $\mu(W_\tau + \tilde{v})$ and $\mu(W_\tau)$. By taking \tilde{v} to be G_S or G_L , we see that the preferences in (3)-(4) can satisfy conditions I and II in the case of immediate gambles so long as there are values of γ and λ for which

$$((W_\tau + 550)^{1-\gamma} + \lambda(W_\tau - 500)^{1-\gamma})^{\frac{1}{1-\gamma}} < (1 + \lambda)^{\frac{1}{1-\gamma}} W_\tau \quad (8)$$

holds for all wealth levels below \$1,000,000, and

$$((W_\tau + 20,000,000)^{1-\gamma} + \lambda(W_\tau - 10,000)^{1-\gamma})^{\frac{1}{1-\gamma}} > (1 + \lambda)^{\frac{1}{1-\gamma}} W_\tau \quad (9)$$

holds for all wealth levels above \$100,000. A quick calculation confirms that $(\gamma, \lambda) = (2, 2)$ satisfies both (8) and (9). Since λ controls sensitivity to losses relative to sensitivity to gains, we need $\lambda > 1.1$ so that the 550/500 bet, with its 1.1 ratio of gain to loss, is rejected.

The intuition for why R-FORA preferences can explain attitudes to G_S and G_L when these gambles are played out immediately, is straightforward. In the case of EU, R-EU, and R-SORA preferences, the difficulty is that the agent is locally risk-neutral, forcing us to push risk aversion over large gambles up to dramatically high levels in order to explain the rejection of G_S , the 550/500 bet. An agent with R-FORA preferences is, by definition, locally risk averse. Risk aversion over large gambles does not, therefore, need to be increased very much to ensure that G_S is rejected.

In the special case where G_S and G_L are played out immediately, then, preferences with first-order risk aversion *can* satisfy conditions I and II. We now show that, in the more realistic and general setting where the gambles are played out with some delay, even preferences with first-order risk aversion have a hard time satisfying these conditions. In particular, while they can easily explain aversion to small, *immediate* gambles, they have great difficulty – in a sense that we make precise – capturing aversion to small, *delayed* gambles. This is a serious concern because, as noted in Section 2, people are just as averse to the 550/500 bet when it is played out with delay as when it is played out immediately.

Before giving an exact statement of the difficulty with R-FORA preferences, we present an informal, static example to illustrate the idea. Consider the following one-period utility function exhibiting first-order risk aversion,

$$w(x) = \begin{cases} x & \text{for } x \geq 0 \\ 2x & \text{for } x < 0 \end{cases} .$$

Such a utility function can easily explain the rejection of an *immediate* 550/500 gamble: an individual with these preferences would assign the gamble a value of $550(\frac{1}{2}) - 2(500)(\frac{1}{2}) = -225$, the negative number signalling that the gamble should be rejected. But how would this individual deal with a small, *delayed* gamble?

In answering this, it is important to recall the essential difference between an immediate and a delayed gamble. The difference is that, while waiting for a delayed gamble's uncertainty to be resolved, the agent is also likely to be exposed to *other* pre-existing risks, such as labor income risk, house price risk, or risk from financial investments. This is not true in the case of an immediate gamble, whose uncertainty, by definition, is resolved at once.

For the R-FORA preferences in (3)-(4), this distinction can have a big impact on whether a gamble is accepted. Suppose that the agent is facing the pre-existing risk $(30,000, \frac{1}{2}; -10,000, \frac{1}{2})$, to be resolved at the end of the period, and is wondering whether to take on an independent, delayed 550/500 gamble, whose uncertainty will also be resolved at the end of the period. The correct way for her to think about this is to merge the new gamble with the pre-existing gamble, and to check whether the combined gamble offers higher utility. Since the combined gamble is

$$(30,550, \frac{1}{4}; 29,500, \frac{1}{4}; -9,450, \frac{1}{4}; -10,500, \frac{1}{4}),$$

the comparison is between

$$30,000(\frac{1}{2}) - 2(10,000)(\frac{1}{2}) = 5,000$$

and

$$30,550(\frac{1}{4}) + 29,500(\frac{1}{4}) - 2(9,450)(\frac{1}{4}) - 2(10,500)(\frac{1}{4}) = 5,037.5.$$

The important point here is that the combined gamble *does* offer higher utility. In other words, the agent would *accept* the small, delayed gamble, even though she would reject an immediate gamble with the same stakes. The intuition is that, since the agent already faces some pre-existing risk, adding a small, independent gamble is a form of diversification, which, even if first-order risk averse, she can enjoy.

This example suggests that, even if the certainty equivalent functional $\mu(\cdot)$ exhibits first-order risk aversion, it may be very difficult to explain the rejection of gambles like G_S , other than in the special case where uncertainty is resolved immediately. In Proposition 1 below, we make the nature of this difficulty precise. In brief, while an individual with R-FORA utility acts in a first-order risk averse manner toward immediate gambles, the presence of pre-existing risks makes her act in a *second*-order risk averse manner toward independent, delayed gambles.

This immediately reintroduces the difficulty noted earlier in our discussion of preferences with second-order risk aversion. Since the agent is second-order risk averse over delayed gambles, and since the delayed gamble G_S is small, she will normally be keen to accept it. In order to explain why it is typically rejected, we need to impose very high local curvature, which, in turn, requires an extreme parameterization. However, such a parameterization usually also implies high global risk aversion and therefore the rejection of large gambles as attractive as G_L . We illustrate this difficulty in Section 3.1 with a formal example.

Proposition 1: Consider an individual with the recursive preferences in (1), where $W(\cdot, \cdot)$ is strictly increasing and twice differentiable with respect to both arguments, and where $\mu(\cdot)$ has the first-order risk averse form in (4).

Suppose that, at time τ , the individual is offered an actuarially favorable gamble $k\tilde{\varepsilon}$ which is to pay off between time τ and $\tau + 1$, and whose payoffs do not affect, and are independent of, time τ information I_τ and future economic uncertainty. Suppose also that, prior to taking the gamble, the distribution of time $\tau + 1$ utility, $\tilde{V}_{\tau+1}$, does not have finite mass at its certainty equivalent $\mu(\tilde{V}_{\tau+1})$.

Then, the individual is second-order risk averse over the new gamble, and, for sufficiently small k , accepts it.

Proof of Proposition 1: See the Appendix.

In simple language, the proposition says that an individual with R-FORA utility is *second-order risk averse* over a delayed, independent gamble, so long as she faces pre-existing risk, an assumption captured here by the statement that “ $\tilde{V}_{\tau+1}$ does not have finite mass at its certainty equivalent $\mu(\tilde{V}_{\tau+1})$.” While Proposition 1 is proven for just one implementation of first-order risk aversion, the argument used in the proof can also be applied, with minor adjustments, to other implementations of first-order risk aversion. For example, by strengthening the assumption that “ $\tilde{V}_{\tau+1}$ does not have finite mass at its certainty equivalent $\mu(\tilde{V}_{\tau+1})$ ” to “ $\tilde{V}_{\tau+1}$ does not have finite mass at any point,” Proposition 1 can be applied when $\mu(\cdot)$ takes the Yaari (1987) rank-dependent expected utility form, which also exhibits first-order risk aversion.⁹

An important step in the proof is an assumption as to how an agent with the preferences in (1) evaluates a delayed gamble \tilde{v} . In their exposition of recursive utility, Epstein and Zin (1989) do not suggest a specific method. We therefore adopt the most natural one, which is that the agent merges the delayed gamble with the other risks she is already taking and checks whether the combination offers higher utility. In other words, she applies the recursive calculation over the time step between τ and $\tau + 1$, and then compares the utility from not taking the gamble,

$$W(C_\tau^*, \mu(\tilde{V}_{\tau+1})) = W(C_\tau^*, \mu(J(\tilde{W}_{\tau+1}))) = W(C_\tau^*, \mu(J((W_\tau - C_\tau^*)\tilde{R}_{\tau+1}^*))), \quad (10)$$

where $J(\cdot)$ is the time $\tau + 1$ value function, $\tilde{R}_{\tau+1}^*$ is the return on invested wealth between τ and $\tau + 1$, and where asterisks denote optimal consumption and portfolio choices; to the

⁹Chapman and Polkovnichenko (2006) note that, if $\tilde{V}_{\tau+1}$ does have finite mass at some point, then rank-dependent expected utility can satisfy conditions I and II. However, for an agent who owns stock or a house, the assumption that $\tilde{V}_{\tau+1}$ does not have finite mass at any point would seem to be a reasonable one.

utility from taking it,

$$W(C_\tau^\dagger, \mu(\tilde{V}_{\tau+1})) = W(C_\tau^\dagger, \mu(J(\tilde{W}_{\tau+1}))) = W(C_\tau^\dagger, \mu(J((W_\tau - C_\tau^\dagger)\tilde{R}_{\tau+1}^\dagger + \tilde{v}))), \quad (11)$$

where optimal consumption and portfolio choices are now denoted by dagger signs. We contrast C_τ^* and $\tilde{R}_{\tau+1}^*$ with C_τ^\dagger and $\tilde{R}_{\tau+1}^\dagger$ as a reminder that, if the agent takes the gamble, her optimal consumption and portfolio choices will be different from what they are when she does not take the gamble.¹⁰

3.1 An Example

We now illustrate the difficulty faced by R-FORA preferences with a formal, intertemporal example. In particular, we show that it is hard for such preferences to explain both the rejection of the delayed gamble G_S and the acceptance of the delayed gamble G_L ; in other words, to satisfy both conditions I and II at the same time.

In our example, we again consider an agent with the R-FORA preferences in (1), (3), and (4). We assume that, initially, the only investment opportunity available to the agent is a risky asset with gross return \tilde{R} between t and $t+1$, where \tilde{R} has the log-normal distribution

$$\log \tilde{R} \sim N(0.04, 0.03), \text{ i.i.d. over time.} \quad (12)$$

In this case, the agent's time t value function takes the form

$$J(W_t) = AW_t \quad (13)$$

for all t , where the constant A is given by

$$\begin{aligned} AW_t &= W(C_t^*, \mu(\tilde{V}_{t+1})) = W(C_t^*, \mu(J(\tilde{W}_{t+1}))) \\ &= W(C_t^*, \mu(J((W_t - C_t^*)\tilde{R}))) = W(C_t^*, A\mu((W_t - C_t^*)\tilde{R})). \end{aligned} \quad (14)$$

Now suppose that, at time τ , the agent is offered a delayed gamble \tilde{v} , where \tilde{v} is either G_S or G_L . As in Proposition 1, \tilde{v} pays off between τ and $\tau+1$, and its payoff does not

¹⁰Strictly speaking, an agent with the preferences in (1), (3), and (4) does not *have* to evaluate the delayed version of G_S by merging it with her pre-existing risk. Since the gamble's uncertainty is resolved at a single instant in the future, she could insert an infinitesimal time interval around that future moment of resolution. Since G_S would be her only source of wealth risk over that interval, her first-order risk aversion would lead her to reject it, consistent with condition I. It is easy, however, to construct a slightly different gamble that is immune to such manipulations. Consider a gamble that, at some point in the future, will deliver a win of \$550 or a loss of \$500 with equal probability. Suppose also that, from now until the final payout, the probability of eventually winning the \$550 is continuously reported. If her pre-existing risk also evolves continuously over time, the agent must necessarily evaluate this 550/500 gamble by merging it with her pre-existing risk, and will therefore accept it. We have found, however, that subjects are as averse to this continuously resolved version of G_S as to the immediate and delayed versions. Once again, the preferences in (3)-(4) have trouble satisfying condition I.

affect, and is independent of, I_t and future economic uncertainty. As a result, whether the agent accepts \tilde{v} or not, her time t value function continues to take the form in equation (13) for all $t \geq \tau + 1$. From (10), the agent’s utility if she does not take \tilde{v} is therefore AW_τ ; from (11), her utility if she does take it is

$$W(C_\tau^\dagger, \mu(J((W_\tau - C_\tau^\dagger)\tilde{R} + \tilde{v}))) = W(C_\tau^\dagger, A\mu((W_\tau - C_\tau^\dagger)\tilde{R} + \tilde{v})). \quad (15)$$

Figure 1 presents the results from comparing AW_τ and (15) when \tilde{v} is G_S or G_L . Our computations set $\beta = 0.9$ and $\rho = -1$, but the results depend little on these choices. The methodology behind the calculations is described in full in the Appendix.

The area marked with “+” signs shows the values of γ and λ for which the agent rejects the delayed 550/500 gamble. Extreme values are required to explain this rejection, with γ exceeding 150 across a wide range of values of λ . The intuition is that, in the presence of the pre-existing risk in (12), the agent acts in a *second*-order risk averse manner towards delayed gambles. In order to explain the rejection of a small, delayed gamble, we therefore need very high local curvature, which, in turn, requires extreme parameters.

The area marked with “x” signs shows the values of γ and λ for which the agent accepts the delayed 20 million/10,000 gamble. There is no overlap between the two shaded regions. In other words, the parameters needed to satisfy condition I are so extreme as to also predict very high global risk aversion, and therefore the rejection of large-stakes gambles as attractive as G_L .

4 Narrow Framing: A Potential Solution

Many utility specifications have difficulty explaining the rejection of a small, independent, actuarially favorable gamble. What kinds of preferences *can* do so? Clearly, first-order risk aversion is a helpful ingredient: it explains why small gambles like G_S are rejected when played out immediately. However, it is not enough: when a first-order risk averse agent evaluates a small, *delayed* gamble, she merges it with her pre-existing risk and, since the resulting diversification is attractive, she happily accepts it. A simple way of explaining the rejection of such a delayed gamble, then, is to suppose that the agent does *not* fully merge it with her pre-existing risk, but that, to some extent, she evaluates it in isolation. By “evaluates it in isolation,” we mean that the agent derives utility directly from the outcome of the gamble, and not just indirectly via its contribution to total wealth, as in traditional models. Equivalently, her utility function depends on the outcome of the gamble over and above what that outcome implies for total wealth risk. This, is “narrow framing.”

We now check that preferences that allow for both first-order risk aversion *and* narrow

framing can easily satisfy conditions I and II, whether the gambles are played out immediately or with delay. A preference specification that incorporates both of these features has recently been developed by Barberis and Huang (2004). In their formulation, time t utility is given by

$$V_t = W \left(C_t, \mu(\tilde{V}_{t+1}) + b_0 E_t \left(\sum_i \bar{v}(\tilde{G}_{i,t+1}) \right) \right), \quad (16)$$

where

$$W(C, y) = ((1 - \beta)C^\rho + \beta y^\rho)^{\frac{1}{\rho}}, \quad 0 \neq \rho < 1, \quad (17)$$

$$\mu(\tilde{V}) = (E(\tilde{V}^{1-\gamma}))^{\frac{1}{1-\gamma}}, \quad 0 < \gamma \neq 1, \quad (18)$$

$$\bar{v}(x) = \begin{cases} x & \text{for } x \geq 0 \\ \lambda x & \text{for } x < 0 \end{cases}, \quad \lambda > 1, \quad (19)$$

and where $\tilde{G}_{i,t+1}$ are specific gambles faced by the agent whose uncertainty will be resolved between time t and $t + 1$.

The term prefixed by b_0 in (16) shows that, relative to the usual recursive specification in (1), utility can now depend on outcomes of gambles $\tilde{G}_{i,t+1}$ over and above what those outcomes mean for total wealth risk: $\tilde{G}_{i,t+1}$ now enters the utility function directly, and not just indirectly via time $t + 1$ utility, \tilde{V}_{t+1} . In other words, the specification in (16) allows for narrow framing, with the parameter b_0 controlling the degree of narrow framing: a b_0 of 0 means no narrow framing at all, while a large b_0 means that $\tilde{G}_{i,t+1}$ is evaluated almost completely in isolation from other risks. First-order risk aversion is also introduced, this time through the piecewise linearity of $\bar{v}(\cdot)$. Since $\bar{v}(\cdot)$ exhibits first-order risk aversion, there is no need for $\mu(\cdot)$ to do so; here, $\mu(\cdot)$ takes a simple power utility form.

Barberis and Huang (2004) propose that an agent with these preferences evaluates an immediate gamble \tilde{x} at time τ by, as before, inserting an infinitesimal time interval $\Delta\tau$ and applying the recursive calculation over this time step. If the gamble is framed narrowly, the utility from taking it is

$$\begin{aligned} W(0, \mu(\tilde{V}_{\tau+\Delta\tau}) + b_0 E_\tau(\bar{v}(\tilde{x}))) &= W(0, \mu(J(\tilde{W}_{\tau+\Delta\tau})) + b_0 E_\tau(\bar{v}(\tilde{x}))) \\ &= W(0, \mu(J(W_\tau + \tilde{x})) + b_0 E_\tau(\bar{v}(\tilde{x}))), \end{aligned} \quad (20)$$

where $J(\cdot)$ is the time $\tau + \Delta\tau$ value function, while the utility from not taking it is

$$W(0, \mu(V_{\tau+\Delta\tau})) = W(0, \mu(J(W_\tau))). \quad (21)$$

A delayed gamble at time τ is evaluated, as before, by applying the recursive calculation over the time step from τ to $\tau + 1$. If the gamble is framed narrowly, the utility from taking it is

$$\begin{aligned} W(C_\tau^\dagger, \mu(\tilde{V}_{\tau+1}) + b_0 E_\tau(\bar{v}(\tilde{x}))) &= W(C_\tau^\dagger, \mu(J(\tilde{W}_{\tau+1})) + b_0 E_\tau(\bar{v}(\tilde{x}))) \\ &= W(C_\tau^\dagger, \mu(J((W_\tau - C_\tau^\dagger)\tilde{R}_{\tau+1}^\dagger + \tilde{x})) + b_0 E_\tau(\bar{v}(\tilde{x}))), \end{aligned} \quad (22)$$

where $J(\cdot)$ is the time $\tau + 1$ value function, $\tilde{R}_{\tau+1}^\dagger$ is the return on invested wealth between τ and $\tau + 1$, and where dagger signs denote optimal consumption and portfolio choices; and the utility from not taking it is

$$W(C_\tau^*, \mu(\tilde{V}_{\tau+1})) = W(C_\tau^*, \mu(J((W_\tau - C_\tau^*)\tilde{R}_{\tau+1}^*))), \quad (23)$$

where optimal consumption and portfolio choices are now denoted by asterisks. We contrast C_τ^* and $\tilde{R}_{\tau+1}^*$ with C_τ^\dagger and $\tilde{R}_{\tau+1}^\dagger$ as a reminder that optimal consumption and portfolio choices change when the gamble is accepted.

Using these expressions, we can check that the preferences in (16)-(19) can indeed satisfy conditions I and II, whether G_S and G_L are immediate or delayed, so long as the gambles are framed narrowly.¹¹ To see the intuition, suppose that the 550/500 bet, whether immediate or delayed, is framed narrowly, so that the agent thinks about it in isolation, to some extent. Since $\bar{v}(\cdot)$ is steeper for losses than for gains, the potential loss of \$500 looms larger than the potential gain of \$550, leading the agent to reject the gamble. In other words, if the agent's first-order risk aversion is focussed directly on the 550/500 bet rather than just on her overall wealth risk, she will be reluctant to take the bet.

We consider the same environment as in the example of Section 3.1. We assume that, initially, the agent's only investment opportunity is a risky asset with gross return \tilde{R} between t and $t + 1$, where \tilde{R} has the distribution

$$\log \tilde{R} \sim N(0.04, 0.03), \text{ i.i.d. over time.} \quad (24)$$

Barberis and Huang (2004) show that, in this case, the agent's time t value function takes the form

$$J(W_t) = AW_t \quad (25)$$

for all t , where the constant A is given by

$$\begin{aligned} AW_t &= W(C_t^*, \mu(\tilde{V}_{t+1})) = W(C_t^*, \mu(J(\tilde{W}_{t+1}))) \\ &= W(C_t^*, \mu(J((W_t - C_t^*)\tilde{R}))) = W(C_t^*, A\mu((W_t - C_t^*)\tilde{R})). \end{aligned} \quad (26)$$

For simplicity, we assume here that the agent does not frame her pre-existing risk narrowly.

Now suppose that, at time τ , the agent is offered a gamble \tilde{v} , where \tilde{v} is either G_S or G_L and is either immediate or delayed. As before, the outcome of \tilde{v} does not affect, and is independent of, I_t and future economic uncertainty. As a result, whether the agent accepts \tilde{v} or not, her time t value function continues to be given by equation (25) for $t \geq \tau + \Delta\tau$ in the case of an immediate gamble, and for $t \geq \tau + 1$ in the case of a delayed gamble. Expressions

¹¹To satisfy conditions I and II, we only need the agent to frame G_S narrowly. For the sake of treating the two gambles symmetrically, we suppose that the agent frames both of them narrowly.

(20)-(23) then allow us to determine the agent's attitudes to immediate and delayed versions of G_S and G_L . The methodological details are described in full in the Appendix.

We set β and ρ , which have little direct influence on attitudes to risk, to 0.9 and -1, respectively. Our calculations then show that, for a range of values of b_0 – including a b_0 as low as 0.001 – it is easy to find parameter pairs (γ, λ) that satisfy conditions I and II, whether the gambles are immediate or delayed. For example, when $b_0 = 0.1$, the pair $(\gamma, \lambda) = (1.5, 3)$ can do so. The intuition is that a λ of 3 generates enough sensitivity to losses to reject the 550/500 bet, with its 1.1 gain to loss ratio, when that bet is evaluated to some extent in isolation; but it does not generate nearly enough sensitivity to losses to reject the highly attractive 20 million/10,000 gamble.

5 Application: Stock Market Participation

We have argued that a wide range of utility functions, including even those with first-order risk aversion, have difficulty explaining an aversion to a small, independent, actuarially favorable gamble while also making sensible predictions about attitudes to large gambles. Utility functions that exhibit both first-order risk aversion *and* narrow framing offer a simple way out of this difficulty.

We now show that our analysis has useful implications for financial markets; for example, for the stock market participation puzzle: the fact that, even though the stock market has a high mean return, many households have historically been reluctant to allocate any money to it (Mankiw and Zeldes 1991, Haliassos and Bertaut 1995).¹²

One approach to the participation puzzle invokes transaction costs of investing in the stock market; another examines whether non-stockholders have background risk that is correlated with the stock market (Heaton and Lucas 1997, 2000, Vissing-Jorgensen 2002). These approaches almost certainly explain some of the observed non-participation, but they may not be able to account for all of it. Polkovnichenko (2005) finds that, even among wealthy households, for whom transaction costs are low, there is still substantial non-participation; and Curcuru, Heaton, Lucas, and Moore (2004) question whether the correlation of the stock market with the background risk of non-stockholders is high enough to explain an equity allocation as low as zero.

¹²Mankiw and Zeldes (1991) report 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 stock. Today, the fraction of households that own stock is closer to 50%. The participation puzzle is therefore primarily the puzzle of why, *historically*, many people did not invest in equities; but also of why, today, people participate more. In this section, we focus on the first part of the puzzle, and return to the second part in Section 6.

Here, we investigate a third approach to the participation puzzle: a *preference*-based approach. In particular, we try to understand what kinds of preferences can generate non-participation in the stock market for parameterizations that we would consider reasonable. We take a reasonable parameterization to be one that makes sensible predictions about attitudes to *large* gambles; for example, one that satisfies condition II, acceptance of the 20 million/10,000 gamble at wealth levels above \$100,000.

It is already understood that preferences that exhibit only second-order risk aversion have difficulty addressing the participation puzzle (Heaton and Lucas 1997, 2000). Our earlier analysis suggests a more surprising prediction: that even preferences with *first*-order risk aversion will have trouble addressing this puzzle. It also suggests that preferences that combine first-order risk aversion with narrow framing – in this case, narrow framing of stock market risk – will have an easier time doing so.

To see the logic behind these predictions, note that, in the absence of narrow framing, an agent must evaluate stock market risk by merging it with her pre-existing risk and checking if the combination is attractive. For most households, stock market risk has a correlation close to zero with other important risks, such as labor income risk, proprietary income risk, and house price risk (Heaton and Lucas 2000). A small equity position is therefore diversifying, and, according to our earlier analysis, the agent *will* find this attractive, even if first-order risk averse. Only an extreme parameterization of her preferences will make her withdraw from the stock market entirely. However, such a parameterization almost always implies high global risk aversion, thereby violating condition II.

A simple way out of this difficulty is to suppose that, when the agent evaluates the stock market, she does not fully merge its risk with the other risks she is already facing, but rather, thinks about it in isolation, to some extent. By focussing the agent’s first-order risk aversion directly on the stock market rather than just on her overall wealth risk, we can generate non-participation more easily.

In making these predictions, we are assuming that our analysis of *independent* gambles is also relevant for gambles that are merely relatively uncorrelated with other risks. While this is likely to be true, the only way to be sure is to check our predictions explicitly in a simple portfolio choice setting. This is what we now do.

Consider an agent who, at the start of each period, has a fixed fraction $\bar{\theta}_N$ of her wealth tied up in a non-financial asset – our so-called pre-existing risk – with gross return

$$\log \tilde{R}_{N,t+1} = g_N + \sigma_N \tilde{\varepsilon}_{N,t+1}, \quad (27)$$

and who is wondering what fraction θ_S of her wealth to invest in the stock market, which has gross return

$$\log \tilde{R}_{S,t+1} = g_S + \sigma_S \tilde{\varepsilon}_{S,t+1}, \quad (28)$$

where

$$\begin{pmatrix} \tilde{\varepsilon}_{N,t} \\ \tilde{\varepsilon}_{S,t} \end{pmatrix} \sim N \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \omega \\ \omega & 1 \end{pmatrix} \right), \text{ i.i.d. over time.} \quad (29)$$

The remaining fraction of her wealth, $1 - \bar{\theta}_N - \theta_S$, is invested in a risk-free asset earning a gross return of R_f , so that the return on total wealth between t and $t + 1$ is

$$\tilde{R}_{W,t+1} = (1 - \bar{\theta}_N - \theta_S)R_f + \bar{\theta}_N \tilde{R}_{N,t+1} + \theta_S \tilde{R}_{S,t+1}. \quad (30)$$

In reality, of course, the fraction of an individual's wealth made up by a non-financial asset like a house is likely to vary over time. Fixing it at $\bar{\theta}_N$ is a simplifying assumption, but is not crucial for our results.

We solve this portfolio problem for three different preference specifications: (a) as a benchmark, the power utility form

$$E_0 \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma}; \quad (31)$$

(b) recursive utility with first-order risk averse certainty equivalent, or R-FORA, as in (1), (3), and (4); and (c), preferences that combine first-order risk aversion with narrow framing of stock market risk, which, following the formulation in (16), can be written as

$$V_t = W \left(C_t, \mu(\tilde{V}_{t+1}) + b_0 E_t(\bar{v}(\tilde{G}_{S,t+1})) \right), \quad (32)$$

where $\tilde{G}_{S,t+1}$, the stock market gamble the agent is taking, is given by

$$\tilde{G}_{S,t+1} = \theta_S (W_t - C_t) (\tilde{R}_{S,t+1} - R_f), \quad (33)$$

and where $W(\cdot, \cdot)$, $\mu(\cdot)$, and $\bar{v}(\cdot)$ are given in (17)-(19).¹³

To match historical annual data, we set the mean g_S and volatility σ_S of log stock market returns to 0.06 and 0.2, respectively. We set the mean g_N and volatility σ_N of log returns on the non-financial asset to 0.04 and 0.03, respectively, and the fraction of non-financial wealth in total wealth, $\bar{\theta}_N$, to 0.75, but our results depend little on the values of these three parameters. A more important parameter is ω , the correlation of stock market risk with the agent's pre-existing risk. Heaton and Lucas (2000) report correlations between stock market risk and three major kinds of pre-existing risk – labor income, proprietary income, and real estate – of -0.07, 0.14, and -0.2 in the aggregate, respectively. They also find that, in the cross-section of households, these correlations rarely exceed 0.2, and, in their simulations,

¹³The simplest way to define the stock market gamble is $\theta_S (W_t - C_t) (\tilde{R}_{S,t+1} - 1)$: the capital allocated to the stock market multiplied by the net return on the stock market. For tractability, we adopt the slight modification of defining the gain or loss on the stock market gamble relative to the risk-free rate R_f . The interpretation is that a stock market return is only considered a gain, and so only delivers positive utility, if it is higher than the return on T-Bills.

consider only correlations in the range $[-0.1, 0.2]$. An ω of 0.1 is therefore a reasonable benchmark value. Finally, we set the gross risk-free rate R_f to 1.02.

For these return process parameters, and for each utility function in turn, we compute the range of preference parameters for which the agent allocates a fraction $\theta_S \leq 0$ of overall wealth to the stock market; in other words, the range of parameters for which, even though the stock market offers a high mean rate of return $-g_S = 0.06$ – she still refuses to participate in it. (The solution technique used for preference specifications (b) and (c) is described in the Appendix.) We then check whether these preference parameters are reasonable; in other words, whether they are consistent with condition II. Specifically, we check whether an agent with wealth of \$100,000, a fraction $\bar{\theta}_N$ of which is invested in the non-financial asset with return distribution in (27), and a fraction $1 - \bar{\theta}_N$ of which is invested in the risk-free asset, would accept a 20 million/10,000 gamble. As in Sections 3.1 and 4, we use $\beta = 0.9$, $\rho = -1$, and $b_0 = 0.1$.

Power utility preferences illustrate the basic puzzle. For the above return process parameters, $\gamma \geq 137$ is required for a $\theta_S \leq 0$ allocation to the stock market; but for such γ , the agent is very risk averse, and would turn down a 20 million/10,000 gamble at a wealth level of \$100,000, violating condition II.

Figures 2 and 3 present results for preferences (b) and (c), respectively. In each figure, the “+” signs indicate the parameters for which the agent chooses a $\theta_S \leq 0$ allocation to the stock market, while the “x” signs show the parameters for which she accepts G_L .

Figure 2 confirms that, for R-FORA preferences, with first-order risk aversion but no narrow framing, it is hard to generate non-participation for reasonable parameter values. In fact, for this implementation of first-order risk aversion, it is impossible: there is no overlap between the two shaded regions. In the absence of narrow framing, a positive position in the stock market is diversifying and therefore desirable, even for a first-order risk averse agent. Extreme parameter values are needed to keep the agent out of the stock market.

The result in Figure 2 may be especially surprising, given that some authors, including Ang, Liu, and Bekaert (2005), appear to show that Gul (1991)-type first-order risk aversion, the type we work with here, *does* address the stock market participation puzzle. In fact, there is no inconsistency. Earlier research has focussed on the special case where the agent has no pre-existing risk, but simply chooses between riskless T-Bills and a risky stock market. In this case, first-order risk aversion over total wealth risk implies first-order risk aversion over stock market risk, and can therefore indeed generate stock market non-participation. Figure 2 shows that, in the more realistic case where the agent has pre-existing risk, this result breaks down: first-order risk aversion can no longer produce non-participation for reasonable parameters.¹⁴

¹⁴Chapman and Polkovnichenko (2006) argue that another implementation of first-order risk aversion,

Figure 3 confirms our second prediction: that, as soon as narrow framing of the stock market is allowed, a wide range of parameter values can deliver a $\theta_S \leq 0$ equity allocation, while still predicting acceptance of G_L , the 20 million/10,000 gamble. For example, the parameter triple $(\gamma, \lambda, b_0) = (1.5, 3, 0.1)$, which satisfies not only condition II but also condition I, predicts non-participation. The intuition is that a λ of 3 generates enough sensitivity to losses to make the stock market appear unattractive when evaluated to some extent in isolation, but not nearly enough to reject G_L .

We have focussed here on the issue of stock market participation. However, the analysis in Sections 3 and 4 applies in the same way to *any* situation where people are reluctant to take on a small, actuarially favorable gamble that is only weakly correlated with pre-existing risks: the tendency, for example, of some stockholders to hold only a small number of stocks, rather than the many stocks recommended for diversification; the refusal of many stock-owning households to diversify their holdings with a position in international equities; or the surprisingly low deductibles many consumers choose in their automobile insurance policies (Cohen and Einav 2005). In each of these cases, first-order risk aversion, on its own, cannot easily explain the facts; but first-order risk aversion combined with narrow framing – of individual stocks, of foreign equity holdings, or of the financial risk of a car accident – can do so more readily.

Barberis and Huang (2006) point out that, in many representative agent models of the stock market, the equity premium is determined by the agent’s attitude, in equilibrium, to taking on an extra dollar of stock market risk that is only weakly correlated with her existing holdings. If these models are to match the high historical equity premium, the agent must be strongly averse to this small amount of weakly correlated risk. The analysis in Section 3 suggests that first-order risk aversion, alone, is unlikely to generate the aversion needed for a high premium, a prediction confirmed by Epstein and Zin (2001) and Barberis and Huang (2006). The analysis in Section 4, however, suggests that a combination of first-order risk aversion and narrow framing – here, narrow framing of stock market risk – will more easily generate the required aversion, and hence also, a high premium. Benartzi and Thaler (1995) and Barberis, Huang, and Santos (2001) confirm that models based on loss aversion over narrowly-framed stock market risk can indeed generate sizeable equity premia.

Yaari’s (1987) rank-dependent expected utility, can generate non-participation more easily. This is only true, however, if the distribution of the agent’s future utility, $\tilde{V}_{\tau+1}$, has finite mass at some point. If $\tilde{V}_{\tau+1}$ does not have finite mass at any point, rank-dependent expected utility encounters the same difficulty as Gul’s (1991) disappointment aversion in explaining non-participation (Haliassos and Hassapis 2001). For an agent who owns a house or a private business, the assumption that $\tilde{V}_{\tau+1}$ does not have finite mass at any point would seem to be a reasonable one. An approach that *can* generate non-participation even for continuously distributed pre-existing risk is one based on ambiguity aversion (Epstein and Schneider 2006). This approach works in much the same way as our own, by inducing first-order risk aversion over the stock market gamble itself.

6 Interpreting Narrow Framing

We have argued that preferences that combine first-order risk aversion with narrow framing may be helpful for understanding attitudes to both independent monetary gambles and the stock market. Of the two features, narrow framing is the more unusual. We therefore end by discussing its interpretation in more detail.

One way narrow framing can arise is if an agent takes non-consumption utility, such as regret, into account. 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 the agent accepts is just one of many risks that she 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, the agent may regret the decision to accept it. Consideration of non-consumption utility therefore leads quite naturally to preferences that depend on the outcomes of gambles over and above what those outcomes mean for total wealth.

Another view of narrow framing is proposed by Kahneman (2003). He suggests that it arises when 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.” When an agent is offered a 50:50 bet to win \$550 or lose \$500, the outcomes of the gamble, \$550 and \$500, are instantly accessible; much less accessible, however, is the distribution of future outcomes the agent faces after integrating the 550/500 bet with all of her other risks. As a result, if the agent thinks about the gamble intuitively, the distribution of the gamble, taken alone, may play a more important role in decision-making than would be predicted by traditional utility functions defined only over wealth or consumption.

By providing the outline of a theory of framing, Kahneman (2003) makes framing a more testable concept than it was before. Thaler, Tversky, Kahneman, and Schwarz (1997) illustrate the kind of test one can do. In an experimental setting, they ask subjects how they would allocate between a risk-free asset and a risky asset over a long time horizon such as 30 years. The key manipulation is that some subjects are shown draws from the distribution of short-term asset returns – the distribution of monthly returns, say – while others are shown draws from a long-term return distribution – the distribution of 30-year returns, say. Since they have the same decision problem, the two groups of subjects should make similar allocation decisions: subjects who see short-term returns should simply use them to infer the more relevant long-term returns. If this requires too much effort, however, Kahneman’s (2003) framework suggests that these subjects will instead use the returns that are most accessible to them, namely the short-term returns they were shown. Since losses occur more often in high frequency data, they will perceive the risky asset to be especially risky and will

allocate less to it. This is exactly Thaler et al.’s (1997) finding.

In Section 5, we addressed the stock market participation puzzle by saying that agents may get utility from the outcome of their stock market investments over and above what that outcome means for their overall wealth; in other words, they may frame the stock market narrowly. Is this a plausible hypothesis?

It seems to us that both the “regret” and “accessibility” interpretations of narrow framing could indeed apply to decisions about the stock market. Allocating some fraction of her wealth to the stock market constitutes a concrete action on the part of the agent – one that she may later regret if her stock market gamble turns out poorly.¹⁵

Alternatively, given our daily exposure, from newspapers, books, and other media, to information about the distribution of stock market risk, such information is very accessible. Much less accessible is any information as to the distribution of future outcomes once stock market risk is merged with the other kinds of risk that people often face. Applying Kahneman’s (2003) framework, judgments about how much to invest in the stock market might therefore be made, to some extent, using a narrow frame. Over time, of course, people may learn that their intuitive decision-making is leading them astray, and may switch to the normatively superior strategy of participating in the stock market. This may be one factor behind the rise in stock market participation over the past 15 years.

The so-called “disposition effect” – the tendency of individual investors to sell stocks in their portfolios that have risen in value since purchase, rather than fallen – suggests that people may not only frame overall stock market risk narrowly, but individual stock risk as well: perhaps the simplest way of explaining the disposition effect is to posit that people receive *direct* utility from realizing a gain or loss on an individual stock that they own.¹⁶

Kumar and Lim (2005) illustrate another approach to testing framing-based theories. They point out that, under the hypothesis that the disposition effect is due to narrow framing, people who are more susceptible to narrow framing will display the disposition effect more. They identify these more susceptible traders as those who tend to execute just one trade, as opposed to several trades, on any given day: if a trader executes just one trade on a given day, the gain or loss for that trade will be more accessible, making narrow framing more likely. The data confirm the prediction: the “one-trade-a-day” traders exhibit the disposition effect more.

¹⁵Of course, investing in T-Bills may also lead to regret if the stock market goes up in the meantime. Regret is thought to be stronger, however, when it stems from having taken an action – for example, moving one’s savings from the default option of a riskless bank account to the stock market – than from not having 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 (Kahneman and Tversky 1982).

¹⁶Odean (1998) argues that other potential explanations of the disposition effect – explanations based on informed trading, taxes, rebalancing, or transaction costs – fail to capture important features of the data.

7 Conclusion

We argue that narrow framing, whereby an agent who is offered a new gamble evaluates that gamble in isolation, separately from other risks she already faces, may be a more important feature of decision-making than previously realized. Our starting point is the evidence that people are often averse to a small, independent gamble, even when the gamble is actuarially favorable. We find that a surprisingly wide range of utility functions, including many non-expected utility specifications, have trouble explaining this evidence; but that this difficulty can be overcome by allowing for narrow framing. Our analysis makes predictions as to what kinds of preferences can most easily address the stock market participation puzzle, as well as other related financial puzzles. We confirm these predictions in a simple portfolio choice setting.

Our analysis does not *prove* that narrow framing is at work in the case of monetary gambles, nor that it is at work in the case of stock market non-participation. However, given the difficulties faced by standard preferences, the narrow framing view may need to be taken more seriously than before. With the emergence of new theories of framing, such as that of Kahneman (2003), we expect to see new tests of framing-based hypotheses. Such tests should, in time, help us learn more about the role that narrow framing plays in individual decision-making.

8 Appendix

Proof of Proposition 1: We prove the proposition for a certainty equivalent functional $\mu(\cdot)$ more general than (4), namely

$$u(\mu(\tilde{V})) = E(u(\tilde{V})) + (\lambda - 1)E((u(\tilde{V}) - u(\mu(\tilde{V})))1(\tilde{V} < \mu(\tilde{V}))), \quad \lambda > 1, \quad (34)$$

where $u(\cdot)$ has a positive first derivative and a negative second derivative. For

$$u(x) = \frac{x^{1-\gamma}}{1-\gamma}, \quad 0 < \gamma \neq 1, \quad (35)$$

this reduces to (4).

Note first that, since $\tilde{V}_{\tau+1}$ does not have finite mass at $\mu(\tilde{V}_{\tau+1})$, a small change in the period $\tau + 1$ value function, $\Delta\tilde{V}_{\tau+1} = \Delta\tilde{V}(\tilde{W}_{\tau+1}, I_{\tau+1})$, changes the certainty equivalent by

$$\Delta\mu = \frac{E(u'(\tilde{V}_{\tau+1})\Delta\tilde{V}_{\tau+1}) + (\lambda - 1)E(u'(\tilde{V}_{\tau+1})1(\tilde{V}_{\tau+1} < \mu)\Delta\tilde{V}_{\tau+1})}{u'(\mu)(1 + (\lambda - 1)\text{Prob}(\tilde{V}_{\tau+1} < \mu))} + o(\|\Delta\tilde{V}_{\tau+1}\|), \quad (36)$$

where μ denotes $\mu(\tilde{V}_{\tau+1})$, $\|x\| = E(|x|)$, and $\lim_{x \rightarrow 0}(o(x)/x) = 0$, by definition.

Denote the gamble $k\tilde{\varepsilon}$ described in the proposition by \tilde{v} . Assume, for now, that the agent does not optimally adjust her time τ consumption and portfolio choice if she takes the gamble. Then,

$$\Delta\tilde{V}_{\tau+1} = J_W(\tilde{W}_{\tau+1}, I_{\tau+1})\tilde{v} + o(\|\tilde{v}\|), \quad (37)$$

which, from (36), implies

$$\Delta\mu = \frac{E(u'(\tilde{V}_{\tau+1})J_W(\tilde{W}_{\tau+1}, I_{\tau+1})\tilde{v}) + (\lambda - 1)E(u'(\tilde{V}_{\tau+1})1(\tilde{V}_{\tau+1} < \mu)J_W(\tilde{W}_{\tau+1}, I_{\tau+1})\tilde{v})}{u'(\mu)(1 + (\lambda - 1)\text{Prob}(\tilde{V}_{\tau+1} < \mu))} + o(\|\tilde{v}\|). \quad (38)$$

Since \tilde{v} is independent of other economic uncertainty, we have

$$\Delta\mu = E(\tilde{v}) \frac{E(u'(\tilde{V}_{\tau+1})J_W(\tilde{W}_{\tau+1}, I_{\tau+1})) + (\lambda - 1)E(u'(\tilde{V}_{\tau+1})1(\tilde{V}_{\tau+1} < \mu)J_W(\tilde{W}_{\tau+1}, I_{\tau+1}))}{u'(\mu)(1 + (\lambda - 1)\text{Prob}(\tilde{V}_{\tau+1} < \mu))} + o(\|\tilde{v}\|), \quad (39)$$

so that, to first order, the certainty equivalent of $\tilde{V}_{\tau+1}$ depends only on the expected value of the gamble \tilde{v} , not on its standard deviation.

To complete the proof, note that the aggregator function $W(\cdot, \cdot)$ does not generate any first-order dependence on the standard deviation of \tilde{v} . In addition, assuming that the agent adjusts her time τ consumption and portfolio choice optimally when accepting \tilde{v} only introduces terms of the second order of \tilde{v} . The agent is therefore second-order risk averse over \tilde{v} .

Illustrative Examples: Acceptance/Rejection of Monetary Gambles

Recursive utility with first-order risk averse certainty equivalent

We now describe the methodology behind the delayed gamble calculations of Section 3.1. An important parameter is the constant A in (13). From (14), A satisfies

$$\begin{aligned} AW_t &= \max_{C_t} W(C_t, A\mu((W_t - C_t)\tilde{R})) \\ &= \max_{C_t} \{(1 - \beta)C_t^\rho + \beta(W_t - C_t)^\rho A^\rho (\mu(\tilde{R}))^\rho\}^{\frac{1}{\rho}} \\ &= \max_{\alpha} W_t \{(1 - \beta)\alpha^\rho + \beta(1 - \alpha)^\rho A^\rho (\mu(\tilde{R}))^\rho\}^{\frac{1}{\rho}}, \end{aligned} \quad (40)$$

where α is the fraction of wealth consumed at time t . The first-order condition is

$$(1 - \beta)\alpha^{\rho-1} = \beta(1 - \alpha)^{\rho-1} A^\rho \mu^\rho, \quad (41)$$

where $\mu = \mu(\tilde{R})$. When substituted into (40), this gives

$$A = (1 - \beta)^{\frac{1}{\rho}} \alpha^{\frac{\rho-1}{\rho}}. \quad (42)$$

Substituting this into (41) leads to

$$\alpha = 1 - \beta^{\frac{1}{1-\rho}} \mu^{\frac{\rho}{1-\rho}}. \quad (43)$$

Therefore, given μ , which can be computed from its definition in (4), we obtain α from (43) and A from (42).

If the agent does not take the delayed gamble \tilde{v} , her utility at time τ is AW_τ . If she does take \tilde{v} , her utility, from (15), is

$$\begin{aligned} &\max_{C_\tau} W(C_\tau, A\mu((W_\tau - C_\tau)\tilde{R} + \tilde{v})) \\ &= \max_{\alpha} W_\tau \left\{ (1 - \beta)\alpha^\rho + \beta(1 - \alpha)^\rho A^\rho \left(\mu(\tilde{R} + \frac{\tilde{v}}{W_\tau(1 - \alpha)}) \right)^\rho \right\}^{\frac{1}{\rho}} \\ &= \hat{A}W_\tau, \end{aligned} \quad (44)$$

where A was computed above. This maximization can be performed numerically. We can then compare \hat{A} to A to see if the agent should take the delayed gamble.

Recursive utility with both first-order risk aversion and narrow framing

We now describe the methodology behind the gamble calculations of Section 4. An important parameter is the constant A in (25). From (26), A satisfies

$$\begin{aligned} AW_t &= \max_{C_t} W(C_t, A\mu((W_t - C_t)\tilde{R})) \\ &= \max_{C_t} \{(1 - \beta)C_t^\rho + \beta(W_t - C_t)^\rho A^\rho E(\tilde{R}^{1-\gamma})^{\frac{\rho}{1-\gamma}}\}^{\frac{1}{\rho}} \\ &= \max_{\alpha} W_t \{(1 - \beta)\alpha^\rho + \beta(1 - \alpha)^\rho A^\rho E(\tilde{R}^{1-\gamma})^{\frac{\rho}{1-\gamma}}\}^{\frac{1}{\rho}}, \end{aligned} \quad (45)$$

where α is the fraction of wealth consumed at time t . The first-order condition is

$$(1 - \beta)\alpha^{\rho-1} = \beta(1 - \alpha)^{\rho-1}A^\rho(E(\tilde{R}^{1-\gamma}))^{\frac{\rho}{1-\gamma}}, \quad (46)$$

which, when substituted into (45), gives

$$A = (1 - \beta)^{\frac{1}{\rho}}\alpha^{\frac{\rho-1}{\rho}}. \quad (47)$$

Substituting this into (46) leads to

$$\alpha = 1 - \beta^{\frac{1}{1-\rho}}(E(\tilde{R}^{1-\gamma}))^{\frac{\rho}{(1-\gamma)(1-\rho)}}. \quad (48)$$

Therefore, we obtain α from (48) and A from (47).

From (20) and (21), the agent takes an immediate gamble \tilde{x} iff

$$A[E(W_\tau + \tilde{x})^{1-\gamma}]^{\frac{1}{1-\gamma}} + b_0E_\tau(\bar{v}(\tilde{x})) > AW_\tau, \quad (49)$$

where A was computed above.

We now turn to the case of delayed gambles. From (23), if the agent does not take a delayed gamble \tilde{x} , her utility is AW_τ , where A was computed above. If she takes the delayed gamble, her utility, from (22), is

$$\begin{aligned} & \max_{C_\tau} W(C_\tau, A\mu((W_\tau - C_\tau)\tilde{R} + \tilde{x}) + b_0E_\tau(\bar{v}(\tilde{x}))) \\ &= \max_{\alpha} W_\tau \left\{ (1 - \beta)\alpha^\rho + \beta(1 - \alpha)^\rho \left[A \left(E\left(\tilde{R} + \frac{\tilde{x}}{W_\tau(1 - \alpha)} \right)^{1-\gamma} \right)^{\frac{1}{1-\gamma}} + b_0E_\tau\left(\bar{v}\left(\frac{\tilde{x}}{W_\tau(1 - \alpha)} \right) \right) \right]^\rho \right\}^{\frac{1}{\rho}}, \\ &= \hat{A}W_\tau, \end{aligned} \quad (50)$$

where A was computed above. This maximization can be performed numerically. We can then compare \hat{A} to A to see if the agent should take the delayed gamble.

Illustrative Examples: Portfolio Choice

We now describe the methodology behind the portfolio choice calculations of Section 5.

Recursive utility with first-order risk averse certainty equivalent

Epstein and Zin (1989) show that, in the i.i.d. setting of Section 5, the consumption-wealth ratio is a constant α , the fraction of wealth allocated to the stock market is a constant θ_S , and the time t value function is $J(W_t) = AW_t$ for all t . The agent's problem becomes

$$\begin{aligned} & \max_{C_t, \theta_S} W(C_t, \mu(\tilde{V}_{t+1})) \\ &= \max_{C_t, \theta_S} W(C_t, A\mu(\tilde{W}_{t+1})) \\ &= \max_{\alpha, \theta_S} W_t [(1 - \beta)\alpha^\rho + \beta(1 - \alpha)^\rho A^\rho (\mu(\tilde{R}_{W,t+1}))^\rho]^{\frac{1}{\rho}}, \end{aligned} \quad (51)$$

where $\tilde{R}_{W,t+1}$ is defined in (30). The consumption and portfolio problems are therefore separable. The portfolio problem is

$$\max_{\theta_S} \mu(\tilde{R}_{W,t+1}), \quad (52)$$

which, given the definition of $\mu(\cdot)$ in (4), can be solved in a straightforward fashion.

Recursive utility with both first-order risk aversion and narrow framing

Barberis and Huang (2004) show that, in the i.i.d. setting of Section 5, the consumption-wealth ratio is a constant α , the fraction of wealth allocated to the stock market is a constant θ_S , and the time t value function is $J(W_t) = AW_t$ for all t . The agent's problem becomes

$$\begin{aligned} AW_t &= \max_{C_t, \theta_S} W(C_t, \mu(\tilde{V}_{t+1}) + b_0 E_t(\bar{v}(\tilde{G}_{S,t+1}))) \\ &= \max_{C_t, \theta_S} W(C_t, A\mu(\tilde{W}_{t+1}) + b_0 E_t(\bar{v}(\tilde{G}_{S,t+1}))) \\ &= \max_{\alpha} W_t[(1 - \beta)\alpha^\rho + \beta(1 - \alpha)^\rho(B^*)^\rho]^{\frac{1}{\rho}}, \end{aligned} \quad (53)$$

where $\tilde{G}_{S,t+1}$ is defined in (33) and

$$B^* = \max_{\theta_S} A[E(\tilde{R}_{W,t+1}^{1-\gamma})]^{\frac{1}{1-\gamma}} + b_0 \theta_S E_t(\bar{v}(\tilde{R}_{S,t+1} - R_f)). \quad (54)$$

The only difficulty with the portfolio problem in (54) is that it depends on the value function constant A . To handle this, note that the first-order condition for consumption in (53) is

$$(1 - \beta)\alpha^{\rho-1} = \beta(1 - \alpha)^{\rho-1}(B^*)^\rho. \quad (55)$$

Substituting this into (53) gives

$$A = (1 - \beta)^{\frac{1}{\rho}} \alpha^{\frac{\rho-1}{\rho}}. \quad (56)$$

The problem can now be solved as follows. Guess a candidate value of α , substitute it into (56) to generate a candidate A , and then solve portfolio problem (54) for that A . Take the B^* that results and use equation (55) to generate a new α . Continue this iteration until convergence occurs. The converged values represent an optimum.

9 References

- Abdellaoui, M., H. Bleichrodt, and C. Paraschiv (2005), "Measuring Loss Aversion: A Parameter-Free Approach," Working paper, Erasmus University.
- Ang, A., G. Bekaert, and J. Liu (2005), "Why Stocks May Disappoint," *Journal of Financial Economics* 76, 471-508.
- 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. and M. Huang (2006), "The Loss Aversion / Narrow Framing Approach to the Equity Premium Puzzle," forthcoming in R. Mehra, ed., *Handbook of Investments: Equity Premium*, Amsterdam: North Holland.
- Barberis, N., M. Huang, and T. Santos (2001), "Prospect Theory and Asset Prices," *Quarterly Journal of Economics* 116, 1-53.
- Barberis, N., M. Huang, and R. Thaler (2003), "Individual Preferences, Monetary Gambles, and the Equity Premium," Working paper 9997, National Bureau of Economic Research.
- Benartzi S., and R. Thaler (1995), "Myopic Loss Aversion and the Equity Premium Puzzle," *Quarterly Journal of Economics* 110, 75-92.
- Chapman, D. and V. Polkovnichenko (2006), "Heterogeneity in Preferences and Asset Market Outcomes," Working paper, Boston College.
- Chetty, R. (2003), "Consumption Commitments, Unemployment Durations, and Local Risk Aversion," Working paper, University of California at Berkeley.
- Cohen, A. and L. Einav (2005), "Estimating Risk Preferences from Deductible Choice," Working paper, Stanford University.
- Curcuro, S., J. Heaton, D. Lucas, and D. Moore (2004), "Heterogeneity and Portfolio Choice: Theory and Evidence," forthcoming in Y. Ait-Sahalia and L. Hansen, eds., *Handbook of Financial Econometrics*, Amsterdam: North Holland.
- Epstein, L. and M. Schneider (2006), "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.
- 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.
- Haliassos, M. and C. Hassapis (2001), "Non-expected Utility, Saving, and Portfolios," *Economic Journal* 111, 69-102.
- 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, 1449-75.
- 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 (1982), "The Psychology of Preferences," *Scientific American* 246, 160-173.
- Kumar, A. and S. Lim (2005), "How Do Decision Frames Influence the Stock Investment Choices of Individual Investors?" Working paper, Notre Dame University.
- Mankiw, N.G. and S. Zeldes (1991), "The Consumption of Stockholders and Non-stockholders," *Journal of Financial Economics* 29, 97-112.
- Odean, T. (1998), "Are Investors Reluctant to Realize their Losses?" *Journal of Finance* 53, 1775-1798.
- Polkovnichenko, V. (2005), "Household Portfolio Diversification: A Case for Rank-dependent Preferences," *Review of Financial Studies* 18, 1467-1501.
- 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., G. Loewenstein, 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.
- Segal, U. and A. Spivak (1990), "First-order versus Second-order Risk Aversion," *Journal of Economic Theory* 51, 111-125.
- Thaler, R., A. Tversky, D. Kahneman, 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 (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.
- Yaari, M. (1987), "The Dual Theory of Choice Under Risk," *Econometrica* 55, 95-115.

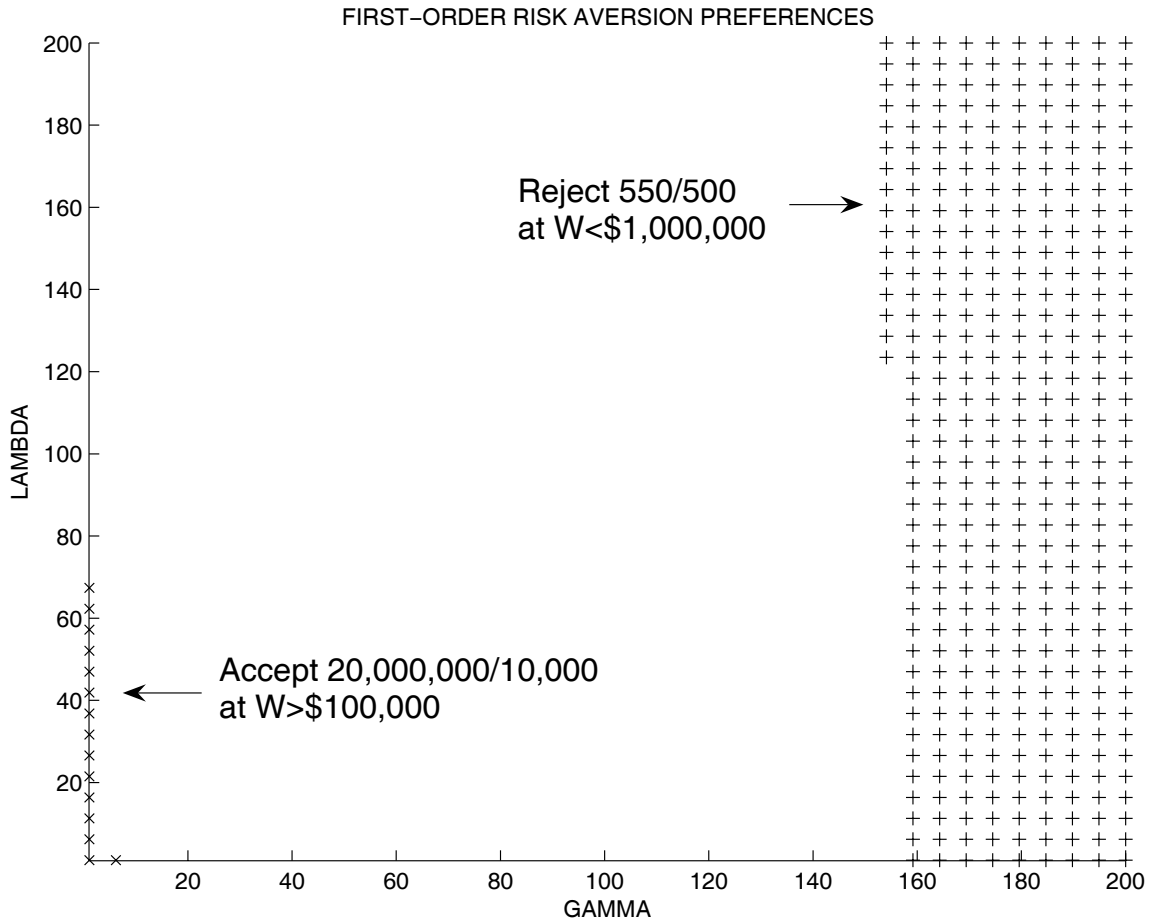


Figure 1. The “+” signs mark the parameter values for which an agent with a recursive utility function with first-order risk averse certainty equivalent *rejects* a 50:50 bet to win \$550 or lose \$500 at a wealth level of \$1,000,000. The “x” signs mark the parameter values for which the agent *accepts* a 50:50 bet to win \$20,000,000 or lose \$10,000 at a wealth level of \$100,000.

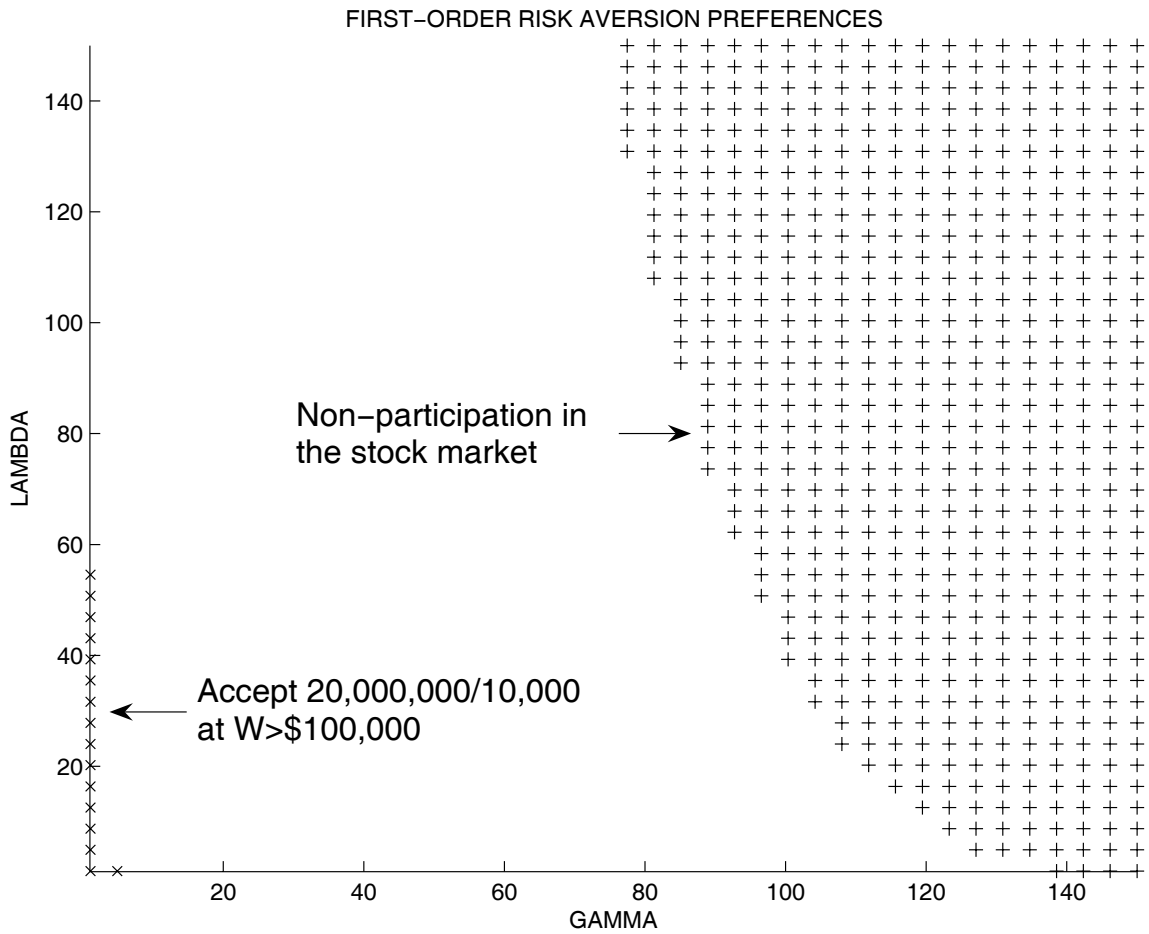


Figure 2. The “+” signs mark the parameter values for which an agent with a recursive utility function with first-order risk averse certainty equivalent does not participate in a stock market offering a high mean return and low correlation with other risks. The “x” signs mark the parameter values for which the agent accepts a 50:50 bet to win \$20,000,000 or lose \$10,000 at a wealth level of \$100,000.

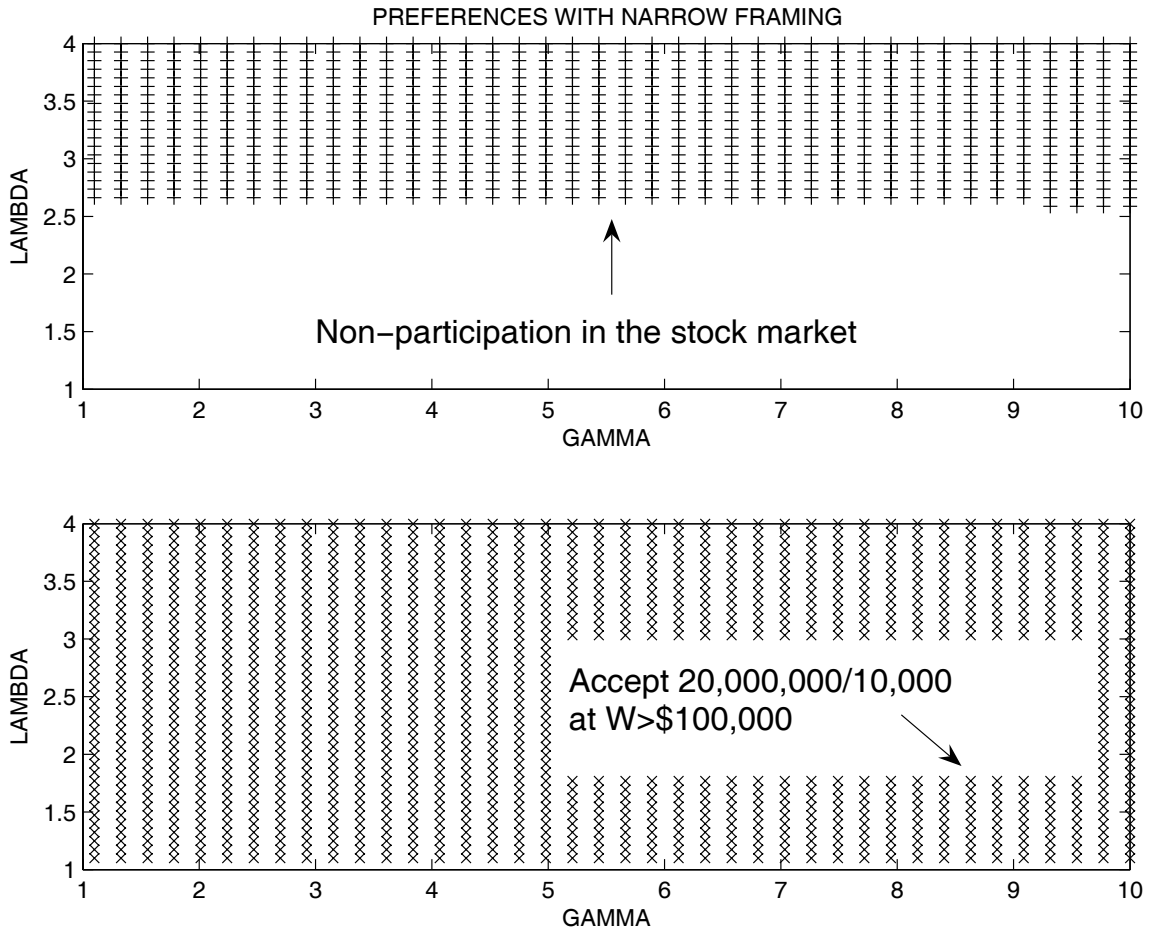


Figure 3. The “+” signs mark the parameter values for which an agent with a recursive utility function that allows for both first-order risk aversion and narrow framing does not participate in a stock market offering a high mean return and low correlation with other risks. The “x” signs mark the parameter values for which the agent accepts a 50:50 bet to win \$20,000,000 or lose \$10,000 at a wealth level of \$100,000.