Federal Reserve Bank of New York
Staff Reports

Anxiety in the Face of Risk

Thomas M. Eisenbach
Martin C. Schmalz

Staff Report No. 610
April 2013
Revised December 2015

Anxiety in the Face of Risk
Thomas M. Eisenbach and Martin C. Schmalz
Federal Reserve Bank of New York Staff Reports, no. 610
April 2013; revised December 2015
JEL classification: D01, D03, D81, G02, G11, G12

Abstract

We model an "anxious" agent as one who is more risk averse with respect to imminent risks than with respect to distant risks. Based on a utility function that captures individual subjects' behavior in experiments, we provide a tractable theory relaxing the restriction of constant risk aversion across horizons and show that it generates rich implications. We first apply the model to insurance markets and explain the high premia for short-horizon insurance. Then, we show that costly delegated portfolio management, investment advice, and withdrawal fees emerge as endogenous features and strategies to cope with dynamic inconsistency in intratemporal risk-return tradeoffs.

Key words: risk premia, insurance, term structure, dynamic inconsistency

---

Eisenbach: Federal Reserve Bank of New York (e-mail: thomas.eisenbach@ny.frb.org). Schmalz: University of Michigan (e-mail: schmalz@umich.edu). This paper was previously distributed under the title "Up Close It Feels Dangerous: Anxiety in the Face of Risk." The authors thank Marianne Andries, Nicholas Barberis, Roland Bénabou, Bruno Biais, Tilman Börgers, Markus Brunnermeier, John Campbell, Sylvain Chassang, Lauren Cohen, Roger Edelen (discussant), Larry Epstein, Xavier Gabaix, Simon Gervais, Christian Gollier, Daniel Gottlieb, Mark Grinblatt, Edoardo Grillo, Andrew Hertzberg (discussant), Burton Hollifield, Daniel Kahneman, Brian Kluger (discussant), Ralph Koijen, David Laibson, Augustin Landier, Stephen Morris, Paulo Natenzon, Juan Ortner, Wolfgang Pesendorfer, Guillaume Plantin, Francesco Sangiorgi (discussant), Norbert Schwarz, Hyun Shin, Marciano Siniscalchi, Satoru Takahashi, Paul Tetlock, and Wei Xiong, as well as audiences at Princeton University, Toulouse School of Economics, the University of Michigan, the University of Vienna, the 2nd Miami Behavioral Finance Conference, the DePaul People and Money Symposium, the 2013 AFA Meetings, the 2013 Finance Down Under Conference, 2014 CICF (Chengdu), and anonymous referees. All remaining errors are those of the authors. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

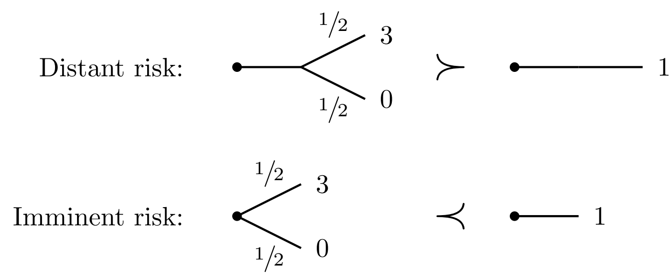

Figure 1: Horizon-dependent risk aversion

1 Introduction

There is ample evidence that people behave in more risk averse ways with respect to risks that are close in time compared to risks that are distant. We term such behavior horizon-dependent risk aversion (HDRA), or more informally 'anxiety'.[1] Despite abundant experimental evidence that people exhibit HDRA preferences, economists have not yet developed a formal way of thinking about such preferences and the implications for economics and finance. This paper takes first steps toward such a framework by modeling an agent whose risk aversion explicitly depends on the temporal distance to the resolution and payoff of a lottery.

Fig. 1 illustrates HDRA with a simple example. In both the top and the bottom comparison the agent has to choose between a risky alternative on the left and a safe alternative on the right. In the top comparison the risk is distant. As a result, the agent has low risk aversion with respect to the gamble. If her risk aversion is low enough, she could choose the risky over the safe alternative. In the bottom comparison the risk is imminent. As a result, the agent has high risk aversion and could choose the safe over the risky alternative. The agent's preference implies different choices depending on the temporal distance of the risk. In particular, she could pull back from risks she previously intended to take, even absent new information and even if her beliefs have not changed for any other reason.

As an intuitive example, consider a parachute jump. An agent could sign up for a jump several days or weeks in advance, thinking the thrill of the jump will be well worth the risk of an accident. However, when looking out the plane's door at the moment of

[1] The New Oxford American Dictionary defines anxiety as a "feeling of worry, nervousness, or unease, typically about an *imminent* event or something with an *uncertain* outcome" (emphasis added). This paper does not discuss *anxiety disorder*, which is a psychopathological condition.

truth, the agent is likely to reconsider and could decide not to jump. Such behavior of parachutists, as well as similar examples, e.g. stage fright of performers, has been studied extensively in the psychology literature (Section 2 provides a discussion). The parachuting example suggests that HDRA has its proximate cause in an emotional reaction to the proximity of risk. We discuss evidence supporting this interpretation in Section 2 as well. However, our analysis does not depend on that interpretation. In our analysis, we postulate an expected utility specification that captures the observed behavior without making a formal claim as to the reasons for such preferences.

The behavior our HDRA preferences capture differs from the behavior captured by related but conceptually orthogonal nonstandard preferences, such as time-varying risk aversion, preference for the timing of resolution of uncertainty, or preferences with non-exponential discounting (which include the quasi-hyperbolic discounting case). Specifically, quasi-hyperbolic discounting represents dynamic inconsistency for intertemporal consumption-savings tradeoffs and gives rise to a demand for illiquid assets and other commitment devices to prevent overconsumption and facilitate saving (Laibson, 1997). Risk is not a central element of such models. In contrast, HDRA represents dynamic inconsistency for intratemporal risk-return tradeoffs and therefore has implications in many domains of decision-making under uncertainty. For example, we show that HDRA can address key features of short-horizon insurance markets which represent puzzles for standard preferences. Moreover, our modeling approach allows us to distinguish between the behavior of 'naive' and 'sophisticated' HDRA agents, with distinct predictions for consumer choice and investor behavior.[2] Such analysis is not possible with preference formulations featuring temptation utilities that imply sophistication throughout (Gul and Pesendorfer, 2001, 2004). Finally, HDRA has the potential to account for features of equilibrium asset prices, as well as particular variation in the cross-sectional pricing of risk, for which nonstandard time preferences have no implications (Luttmer and Mariotti, 2003).

Modeling preferences with HDRA presents several challenges, particularly if one wants to maintain dynamic consistency for intertemporal tradeoffs. We show that in a time-separable framework with more than two periods, HDRA necessarily leads to dynamic inconsistency in consumption even when the increased flexibility of nonex-

[2]For a related analysis with dynamically inconsistent time preferences see, e.g. O'Donoghue and Rabin (1999).

ponential discounting is taken into account. This insight complements that of Strotz (1955): We show that to achieve dynamic consistency, not only does discounting have to be exponential, we show that it is also necessary that the utility indexes be identical. The only way to maintain time-separability and have HDRA without dynamic inconsistency for intertemporal consumption tradeoffs is to restrict analysis to a two-period setting. This is the approach we choose in this paper because it allows for analytical transparency and application to a wide range of settings. We drop time separability in Andries, Eisenbach, and Schmalz (2014) and develop generalized Epstein and Zin (1989) preferences to derive asset pricing implications in a fully dynamic model.

After discussing these modeling challenges, we first apply our model to consumer demand for insurance and for commitment devices to take risk. Given the potential of dynamically inconsistent risk-taking, we can distinguish between 'naive' and 'sophisticated' HDRA agents. Only naive agents will buy very high-priced short-term insurance in the presence of cheaper alternatives that, however, would require more foresight. By incorporating such decisions in an otherwise standard framework, we show that HDRA behavior is not necessarily inconsistent with standard von Neumann-Morgenstern utility functions, resolving the puzzle posed by Eisner and Strotz (1961).

By contrast, only an agent who is sophisticated about her dynamic inconsistency is willing to pay for commitment devices to take risk. In particular, lacking the resolve to personally manage an equity portfolio, she is willing to pay a fee to delegate her investment decisions. Thus, sophisticated agents with HDRA preferences generate a demand for delegated portfolio management, even if these services are costly and known to underperform passive benchmarks that are available at low costs (Gruber, 1996). Moreover, the HDRA model predicts that demand for investment advice is particularly strong for agents who would otherwise not invest in risky assets at all, as documented by Foerster et al. (2014). Our results therefore suggest that firms respond to the presence of HDRA agents in the population and that features of different markets can be understood by allowing for heterogeneity in agents' levels of sophistication.

Finally, we show in a stylized setting that investors with HDRA require more compensation for short-run risks than for long-run risks. This result is suggestive of the downward-sloping term structure of risk premia in equity markets first documented by van Binsbergen, Brandt, and Koijen (2012) and van Binsbergen et al. (2013) and found also in housing and derivatives markets (Giglio et al., 2013; Dew-Becker et al., 2014;

Andries et al., 2015). A proper treatment of this phenomenon requires incorporating HDRA into a more specialized asset-pricing framework as in Andries et al. (2014). However, the two-period model developed in the present paper conveys the intuition why HDRA leads to a downward-sloping term structure, whereas standard asset pricing models predict a flat or upward-sloping term structure.

Many more applications are left to future research. One example is the relationship between HDRA preferences and endogenous information acquisition and belief formation. We show in Eisenbach and Schmalz (2015) why it can be beneficial to a sophisticated agent with HDRA to hold overly precise beliefs ("overconfidence") and how such rational self-delusion can lead to excessive risk taking.

The paper proceeds as follows. Section 2 presents experimental evidence for our main assumption – risk aversion decreases with temporal distance – and discusses potential origins of HDRA preferences. Section 3 discusses challenges in modeling HDRA and derives the model used in this paper. We apply the model in Section 4 to analyze the implications of HDRA in insurance markets. In Section 5, we investigate how sophisticated HDRA agents respond to the potential of dynamically inconsistent behavior with respect to intratemporal risk-return tradeoffs. Section 6 sketches out an application to asset pricing. Section 7 contrasts HDRA with related nonstandard preferences. Section 8 concludes.

2 Experimental evidence and potential origins

Horizon-dependent risk aversion is very well documented experimentally. Subjects tend to be more risk averse when a risk is temporally close than when it is distant, in both across-subject and within-subject studies. In this section, we first review the experimental evidence that temporal distance affects risk-taking behavior and then discuss potential origins of this phenomenon.

Jones and Johnson (1973) have subjects participate in a simulated medical trial for a new drug; each subject has to decide on the dose of the drug to be administered. The subjects are told that the probability of experiencing unpleasant side effects increases with dosage – but so does monetary compensation. More risk averse subjects should then choose lower doses than less risk averse subjects. The study finds that subjects choose higher doses when the drug is to be administered the next day than when it is to be administered immediately. Interestingly, the difference disappears if the decision can be revisited the next day (no commitment), suggesting that some subjects may anticipate their preference reversals.

Welch (1999) documents preference reversals caused by stage fright. Of experiment participants who agree to tell a joke in front of a class the following week in exchange for \$1, Welch finds that 67% "chicken out" when the moment of truth arrives. In contrast, none of those who decline initially change their minds.

Noussair and Wu (2006) as well as Coble and Lusk (2010) use the protocol of Holt and Laury (2002), a widely used method in experimental economics, to elicit risk aversion. Subjects are presented with a list of choices between two binary lotteries. The first lottery always has two intermediate prizes, e.g. $(\$10.00, \$8.00)$, while the second lottery always has a high and a low prize, e.g. $(\$19.25, \$0.50)$. Going down the list, only the respective probabilities of the two prizes change, varying from $(0.1, 0.9)$ to $(0.9, 0.1)$. As probability mass shifts from the second prize to the first prize of both lotteries, the second lottery becomes increasingly attractive compared to the first lottery. Subjects are asked to pick one of two lotteries for each of the probability distributions. The probability distribution at which a subject switches from the "safe" lottery to the "risky" lottery is a proxy for the subject's risk aversion. Noussair and Wu (2006) use this protocol for a within-subject design with real payoffs, having each subject make choices for resolution and payout to occur immediately and also for risks and payouts that occur three months later. The study finds that the number of subjects higher risk aversion for the present than for the future is four times higher than the number of subjects with lower risk aversion for the present. Coble and Lusk (2010) use the Holt-Laury protocol for an across-subject design and find the same pattern, with average risk aversion decreasing as the temporal distance of the risk increases.

In a different type of experiment, Baucells and Heukamp (2010) let subjects choose between two binary lotteries, a "safer" one and a "riskier" one. Different treatments vary the delay between when subjects choose and when the lotteries are resolved and paid out. The study finds that more subjects choose the riskier lottery as the delay increases. Sagristano, Trope, and Liberman (2002) also have subjects choose between two lotteries that differ in risk and find the same effect of temporal horizon.

Finally, some studies elicit risk aversion by asking subjects for their certainty equivalents for different lotteries; a lower certainty equivalent corresponds to higher risk aversion. In Onculer (2000), subjects state their certainty equivalents for a lottery to be resolved and paid immediately, as well as for the same lottery to be resolved and paid in the future. The study finds that subjects state significantly lower certainty equivalents for the immediate lottery than for the future lottery. Abdellaoui, Diecidue, and Onculer (2011) conduct a similar study with real payoffs and find equivalent results.

While our model uses an expected utility formulation with standard intertemporal lotteries over monetary outcomes, we find it intuitively plausible that HDRA arises due to the effect of emotions on decision making and the fact that emotional responses are stronger for more salient cues.

Loewenstein, Weber, Hsee, and Welch (2001) point out that cognitive evaluations of risk do not depend on temporal distance; in contrast, emotional reactions to risk such as fear and anxiety increase as the risk draws closer (see also Loewenstein, 1987, 1996; Monat and Lazarus, 1991; Paterson and Neufeld, 1987). The authors point out that when such departures between thoughts and emotions occur, feelings often exert a dominating influence on behavior. As a result, agents tend to behave in more risk averse ways with respect to risks at shorter horizons, even when cognitive evaluations of the risk remain constant.

Indeed, research in psychology documents a robust link between temporal proximity of risk and 'anxiety' as an emotional response. Some studies even show both horizon-dependent risk aversion preferences and an 'anxiety-prone' emotional response jointly. For example, the study by Jones and Johnson (1973), previously discussed in Section 2, also measures higher stress levels for subjects deciding over immediate doses than for subjects deciding over delayed doses. Monat (1976) and Breznitz (2011) inform subjects that they will receive an electric shock (presumably of an uncertain strength given a subject-specific scale). The temporal distance varies across different treatment groups. Heart rate, and in the latter study also galvanic skin response and self-reported anxiety, are all higher when the shock is closer in time. Fenz and Epstein (1967), Fenz and Jones (1972), and Roth, Breivik, Jørgensen, and Hofmann (1996) investigate the emotional response of parachutists approaching the time of a jump. Novice parachutists exhibit a similar dynamic of physiological measures and self-reports of anxiety as in the above experiments, while expert parachutists have a somewhat attenuated response to the proximity of the jump, suggesting an adaptive nature of 'anxiety.' Lo and Repin (2002) and Lo, Repin, and Steenbarger (2005) find similar psychophysiological responses to risk taking among securities traders.

While intuitively plausible, we do not claim that emotions are indeed the driver of the observed behavior. One reason is that in some theories of emotion, cognition drives emotion rather than the other way around (Gross and Barrett, 2011). As a result, cognition and emotions cannot be as cleanly separated as suggested above.

Trope and Liberman (2003) offer "construal theory," widely accepted by psychologists, to explain choice behavior that differs by horizon. The theory proposes that the mental representation of events depends on the temporal distance to the event. Indeed, neurological evidence indicates that "separate neural systems value immediate and delayed monetary rewards" (McClure, Laibson, Loewenstein, and Cohen, 2004). As a consequence of different representations that come in different levels of abstraction, people make different decisions. The objective of this paper is to provide an economic model of such behavior that is useful to study the implications of these notions from the psychology literature.

In sum, while the evidence on horizon dependence in the emotional response to risk as well as theories used by psychologists to explain horizon-dependent decision making are consistent with and plausibly linked to the horizon-dependent risk choices we discuss, we make no claim as to HDRA's emotional or psychological origins. We take the standpoint of traditional economics: we observe choice and infer preferences, which we subsequently take as given when modeling behavior in different contexts. In the following sections, we examine a preference that reflects the experimental choice behavior without relying on any specific underlying driver of such behavior.

3 Model

Capturing the experimental evidence on horizon-dependent risk aversion discussed in Section 2 raises several questions. On the one hand we need to decide on how to model the observed preferences. As we show below, building on the time-separable utility framework widely used in economics presents important challenges. On the other hand, due to the potential for dynamic inconsistency, we need to decide how to solve the model: We can assume the agent to be naive or sophisticated about the dynamic inconsistency.

Suppose we want to build on the standard time-separable model of expected utility. Denoting an uncertain intertemporal consumption stream from period t onwards by

$\tilde{C}_t = (\tilde{c}_t, \tilde{c}_{t+1}, \tilde{c}_{t+2}, \ldots, \tilde{c}_T)$, we can generalize the standard model by using the following utility function:

$$U_t(\tilde{C}_t) = \mathbb{E}\big[\delta_0 u_0(\tilde{c}_t) + \delta_1 u_1(\tilde{c}_{t+1}) + \cdots + \delta_{T-t} u_{T-t}(\tilde{c}_T)\big]. \tag{1}$$

This utility function has both general discount factors δ_h and general von Neumann-Morgenstern utility indexes u_h for every horizon h relative to the current period t. In the standard model, discounting is geometric, $\delta_h = (\delta)^h$ for all h, and the utility index do not depend on the horizon, $u_h = u$ for all h.

Consider two lotteries $\tilde{x}$ and $\tilde{y}$ such that $\tilde{x} = \tilde{y} + \mu + \tilde{\varepsilon}$ with μ a constant and $\tilde{\varepsilon}$ a mean-zero lottery independent of $\tilde{y}$; such lotteries represent a typical risk-reward tradeoff in which $\tilde{x}$ is "high risk, high reward" and $\tilde{y}$ is "low risk, low reward." To capture the evidence of an HDRA agent in period t who prefers the risky lottery $\tilde{x}$ if it is delayed, e.g. to period $t+1$, but prefers the safe lottery $\tilde{y}$ if it is immediate, the general framework (1) has to satisfy:

$$\text{For } h = 1: \quad \mathbb{E}\big[\delta_1 u_1(\tilde{x})\big] > \mathbb{E}\big[\delta_1 u_1(\tilde{y})\big] \tag{2}$$

$$\text{For } h = 0: \quad \mathbb{E}\big[\delta_0 u_0(\tilde{x})\big] < \mathbb{E}\big[\delta_0 u_0(\tilde{y})\big]. \tag{3}$$

Given the assumptions on $\tilde{x}$ and $\tilde{y}$, these conditions can be satisfied only if u_0 is more risk averse than u_1. It is important to note that the discount factors δ_0 and δ_1 cancel out of the two conditions above. This illustrates the conceptual difference between intratemporal risk tradeoffs and intertemporal consumption tradeoffs; the discount factors of a time-separable model affect the latter but not the former. The experimental evidence can therefore *not* be addressed by relaxing the standard assumption of geometric discounting.

The experimental evidence cited in Section 2 mainly contrasts imminent risks with delayed risks rather than risks with differing delays. We can therefore simplify the general model in (1) using only two utility indexes v and u by setting $u_0 = v$ for immediate risks and $u_h = u$ for delayed risks at all horizons $h \geq 1$ and by assuming that v is more risk averse than u by the Arrow-Pratt measure of absolute risk aversion:

$$U_t(\tilde{C}_t) = \mathbb{E}\big[\delta_0 v(\tilde{c}_t) + \delta_1 u(\tilde{c}_{t+1}) + \cdots + \delta_{T-t} u(\tilde{c}_T)\big] \tag{4}$$

$$\text{with} \quad -\frac{v''(c)}{v'(c)} \geq -\frac{u''(c)}{u'(c)} \quad \text{for all } c. \tag{5}$$

Given that the phenomenon of HDRA is conceptually orthogonal to phenomena of horizon-dependent impatience such as Laibson (1997) (see the discussion in Section 7), it would be desirable to keep the model free of any impatience elements so that we can cleanly identify the implications of HDRA. However, the time-separable approach taken in (1) and (4) has the problem of confounding the dynamically inconsistent risk preferences with dynamically inconsistent time preferences. Consider the following two deterministic consumption streams:

$$C_t = (c, c_L, c, c, \ldots) \quad \text{and} \quad C'_t = (c, c, c_H, c, \ldots) \quad \text{with} \quad c_L < c_H. \tag{6}$$

The two consumption streams only differ in periods $t+1$ and $t+2$ and choosing between the two involves the intertemporal tradeoff whether to receive the smaller c_L earlier or the larger c_H later. Since the consumption streams are deterministic, we want the HDRA agent to evaluate them the same in period t and in period $t+1$. This imposes a restriction on the utility function (4):

$$U_t(C_t) = U_t(C'_t) \quad \Leftrightarrow \quad U_{t+1}(C_{t+1}) = U_{t+1}(C'_{t+1}). \tag{7}$$

First, note that $U_t(C_t) = U_t(C'_t)$ implies:

$$\delta_1 u(c_L) + \delta_2 u(c) = \delta_1 u(c) + \delta_2 u(c_H)\,. \tag{8}$$

Second, note that $U_{t+1}(C_{t+1}) = U_{t+1}(C'_{t+1})$ implies:

$$\delta_0 v(c_L) + \delta_1 u(c) = \delta_0 v(c) + \delta_1 u(c_H)\,. \tag{9}$$

Combining Eq. (8) and Eq. (9) we get:

$$\frac{v(c_L) - v(c)}{u(c_L) - u(c)} = \frac{(\delta_1)^2}{\delta_0 \delta_2}. \tag{10}$$

We want this to hold for arbitrary c_L and c, which implies:

$$\frac{v'(c)}{u'(c)} = \frac{(\delta_1)^2}{\delta_0\delta_2} \quad \text{for all } c. \tag{11}$$

For any general horizon-dependent discounting, $(\delta_1)^2/(\delta_0\delta_2)$ is always a constant so to satisfy (11) the utility indexes v and u can only differ by an affine transformation. This, however, rules out that v and u have different levels of risk aversion as required to represent HDRA behavior. Notice that an implication is that flexible time-discount factors cannot be used to render a utility function with horizon-dependent risk aversion dynamically consistent. In other words, to attain dynamic consistency in a time-separable model, not only do time-discount factors need to be exponential (as noted by Strotz, 1955), we point out that utility indexes also have to be identical. As a consequence, the assumption of identical utility indexes ubiquitous in the literature is a special case with considerable loss of generality.

The key problem revealed in the above discussion is the link between intertemporal substitution and risk aversion inherent in the time-separable model of (1) or (4). There are two solutions to this problem. The first is to depart from the time-separable model to a model that separates time and risk preferences – in the spirit of Epstein and Zin (1989) – yet allows for risk preferences to depend on the horizon and be dynamically inconsistent.[3] The agent's optimization problem then involves a game against her future selves and involves significant analytical complexity. We follow this approach in Andries et al. (2014).

In this paper, we use the second solution to the problem which is much simpler: We restrict analysis to a two-period model with $t = 0, 1$. As the example with the lotteries $\tilde{x}$ and $\tilde{y}$ above illustrates, a two-period setting is sufficient to represent the behavior revealed by the experimental evidence on HDRA and it generates a rich set of implications discussed in the following sections. In contrast to a setting with more than two periods, however, there is no scope for dynamically inconsistent time preferences. Dynamic inconsistency in intertemporal tradeoffs requires two periods $t = 1, 2$ for the tradeoff and at least one prior period $t = 0$ in which the agent resolves the tradeoff differently than in $t = 1$. The restriction to a two-period setting allows us to cleanly

[3]Note that recursive utility formulations such as those used in Epstein and Zin (1989) are also limited to the special case of dynamic consistency, by construction.

identify implications of HDRA without having to worry about confounding influences of conceptually orthogonal theories, and especially quasi-hyperbolic discounting. We therefore use the following setup in this paper:

$$U_0(\tilde{c}_0, \tilde{c}_1) = \mathbb{E}\big[v(\tilde{c}_0) + \delta u(\tilde{c}_1)\big] \tag{12}$$

$$\text{and} \quad U_1(\tilde{c}_1) = \mathbb{E}\big[v(\tilde{c}_1)\big] \tag{13}$$

$$\text{with} \quad -\frac{v''(c)}{v'(c)} \geq -\frac{u''(c)}{u'(c)} \quad \text{for all } c. \tag{14}$$

Finally, the potential for dynamic inconsistency raises the question of how to solve the model. The agent can be naive and not realize that in the future she will not want to follow through with plans made in the present. In that case, the agent simply maximizes the utility function U_t in every period t, choosing the optimal values for the future from the perspective of period t and (wrongly) assuming that she will not re-optimize and choose different values in the future. Alternatively, the agent can be sophisticated in the tradition of Strotz (1955) and optimize subject to the constraint that she will also optimize in the future. We will consider the implications of both naive and sophisticated behavior in the following analysis. The dynamic inconsistency also raises issues for welfare analysis. Since there is no generally accepted welfare criterion for dynamically inconsistent agents, we focus on purely positive analysis in this paper.

4 Short-horizon insurance markets

The prediction that individuals are willing to pay high premia for short-horizon insurance finds empirical support in insurance markets. An example is accidental-death insurance in the context of commercial airline flights. Sold at airports just before the flight they cover, these policies were very popular in the U.S. between the 1950s and 1970s and still are in countries such as Japan (known as *yokouhoken-jidouhanbaiki*, 旅行保険自動販売機).[4] MacKinlay (1963) reports that a single group of underwriters for air trip insurance in the US in the early 1960s sold a notional amount of $84.6 billion ($650.7 billion in 2015 dollars) of insurance, collecting premiums of $3.4 million and

[4]The U.S. market for this insurance dried up because of legal issues related to selling insurance with vending machines (see, e.g., Steven v. Fidelity & Casualty Co. (1962) 58 Cal.2d 862).

paying out \$1.4 million in losses, thus generating a profit margin of 59%.[5]

In a static setting, Eisner and Strotz (1961) show that additional life insurance tied to specific events can violate von Neumann-Morgenstern preferences. To paraphrase their setup, let A denote the agent's wealth and p_0 the probability of death due to any cause, which reduces the wealth available to the agent's dependents by L, e.g. because of lost earnings. Assuming that the agent values the wealth available to her and her dependents by the same von Neumann-Morgenstern utility u, she should buy an amount I_0 of life insurance priced at $\pi_0 I_0$ such that

$$\frac{u'(A - \pi_0 I_0^* - L + I_0^*)}{u'(A - \pi_0 I_0^*)} = \frac{(1 - p_0)\,\pi_0}{p_0\,(1 - \pi_0)} =: \mu_0, \tag{15}$$

i.e. the ratio of marginal utilities equals a "markup" μ_0.

Eisner and Strotz (1961) then assume that after buying the optimal amount of general life insurance the agent unexpectedly needs to travel by air and is offered additional coverage in case of death due to a plane crash. Let p_1 denote the probability of death due to a plane crash and I_1 the amount of additional insurance priced at $\pi_1 I_1$. We assume without loss of generality that the travel decision doesn't affect the overall probability of death p_0. Given the existing insurance of I_0^*, the agent solves:

$$\begin{aligned}\max_{I_1 \geq 0} \Big\{ & p_1 u(A - \pi_0 I_0^* - \pi_1 I_1 - L + I_0^* + I_1) \\ & + (p_0 - p_1)\, u(A - \pi_0 I_0^* - \pi_1 I_1 - L + I_0^*) \\ & + (1 - p_0)\, u(A - \pi_0 I_0^* - \pi_1 I_1) \Big\}\end{aligned} \tag{16}$$

Evaluating the derivative with respect to I_1 at $I_1 = 0$, yields that the agent should buy additional insurance, $I_1^* > 0$, if and only if

$$\frac{u'(A - \pi_0 I_0^* - L + I_0^*)}{\frac{p_0 - p_1}{1 - p_1} u'(A - \pi_0 I_0^* - L + I_0^*) + \frac{1 - p_0}{1 - p_1} u'(A - \pi_0 I_0^*)} > \frac{(1 - p_1)\,\pi_1}{p_1\,(1 - \pi_1)} =: \mu_1, \tag{17}$$

i.e. a modified ratio of marginal utilities exceeds the markup μ_1 of the additional insur-

[5] Another example of insurance apparently targeted at HDRA agents are additional warranties offered to customers in retail settings. For electronics and appliances, 75% of customers buy such additional warranties (Desai and Padmanabhan, 2004) which offer retailers margins of 44–77% (Maronick, 2007) and can account for up to 50% of a retailer's profit (Baird and Benady, 1996).

ance. This leads to the following result on non-HDRA agents – the only type considered by Eisner and Strotz (1961). (All proofs are relegated to Appendix A.)

Proposition 1 (Adapted from Eisner and Strotz, 1961)**.** *If unexpected additional life insurance has a higher markup than existing general-coverage life insurance, $\mu_1 > \mu_0$, then a non-HDRA agent should not buy it.*

Eisner and Strotz (1961) therefore view the popularity of such insurance policies in light of their high markups as a puzzle. However, we can resolve the puzzle by taking into account the possibility of HDRA agents and the timing of the purchase decision relative to the risk covered.

Suppose an HDRA agent with long-term risk aversion u has already purchased long-term general life insurance I_0^* given by (15). As in Eisner and Strotz (1961), the agent is then unexpectedly offered additional insurance for a risk at short horizon which she evaluates with the more risk averse utility v. The agent finds it optimal to purchase additional insurance if and only if

$$\frac{v'(A - \pi_0 I_0^* - L + I_0^*)}{\frac{p_0 - p_1}{1 - p_1} v'(A - \pi_0 I_0^* - L + I_0^*) + \frac{1 - p_0}{1 - p_1} v'(A - \pi_0 I_0^*)} > \mu_1, \tag{18}$$

i.e. the analog of condition (17) but using v instead of u. Since v is more risk averse than u and I_0^* is chosen by u, we have

$$\frac{v'(A - \pi_0 I_0^* - L + I_0^*)}{v'(A - \pi_0 I_0^*)} > \frac{u'(A - \pi_0 I_0^* - L + I_0^*)}{u'(A - \pi_0 I_0^*)} \tag{19}$$

$$= \mu_0 \quad \text{by the first-order condition (15)}, \tag{20}$$

which implies the following result.

Proposition 2. *If v is sufficiently more risk averse than u, an agent with HDRA preferences (v, u) finds it optimal to buy unexpected additional life insurance even if it is relatively more expensive than her pre-existing policy, $\mu_1 > \mu_0$.*

As Krantz et al. (2007) note, even if more comprehensive insurance plans exist and are cheaper, using them instead of the ones offered at the airport requires forethought as well as sophistication about one's own behavior in the face of risk. In the context of our model, agents who buy additional insurance at the airport are either faced with the option unexpectedly or are naive about their dynamically inconsistent risk preferences. Of course, the degree of sophistication about dynamic inconsistency can vary across individuals (DellaVigna and Malmendier, 2006) and even across contexts (Laibson, 2015). Hence, it is not surprising that firms adapt their offerings to target particular subgroups of the population that are characterized by different levels of sophistication about their HDRA.

5 Commitment devices and institutional responses

An agent who plans for tomorrow according to preference u, but realizes that her future self will disagree with these plans (because she will have preference v), could try to find ways to commit to a future plan of action. While Schelling (1984) and others have discussed the ethical aspects that such a possibility brings about, the present discussion is only concerned with the fact that – and the question how – the agent can restrict her future self's behavior, simply by virtue of having a first-mover advantage over her future selves. Similarly, we make no claim as to the normative implications (Gollier, 2012) but merely provide a positive model.

Hiring an agent to carry out future risk-taking decisions according to the current self's preferences is one way to prevent future actions from conflicting with the current self's plans. In an investment setting, it can be the case that the anxious self is too risk averse to invest in equity, although the agent realizes that doing so has long-run benefits compared to saving in a less risky alternative that yields lower average returns. In this situation it makes sense for the agent to delegate investment decisions to a portfolio manager. In fact, Vanguard explicitly lists behavioral coaching as one of the benefits of its investment advisory services, stating that "some of the most significant opportunities [for the advisor] to add value occur [...] when clients are tempted to abandon their well-thought-out investment plan" (Kinniry, Jaconetti, DiJoseph, and Zilbering, 2014). They posit that behavioral coaching generates an "advisor's alpha" of 150 basis points, based on a Vanguard study comparing the average returns of IRA account holders who make unadvised changes to their portfolio allocation to the average returns of corresponding target-date retirement funds (Weber, 2013). These considerations suggest that the value of delegated portfolio management, in contrast to arrived notions, may not lie in picking underpriced stocks. Rather, delegated portfolio management can help individuals prevent what Campbell (2006) identifies as the number one investment mistake households make: nonparticipation in risky asset markets. In other words, the appropriate benchmark to evaluate the performance of delegated portfolio and investment advice for HDRA agents is not a passive risky asset return, but the risk-free rate.

The following model formalizes this intuition in a setting with two periods, $t = 0, 1$. Going backwards, at the beginning of period 1, the HDRA agent has to form a portfolio (ϕ_1, ξ_1) consisting of a risky asset and a risk-free asset. The price of the risky asset is p and it pays off a random $\tilde{x}$ at the end of period 1. In period 0, the agent decides whether to delegate the investment decision to a manager. The manager charges a fee $f > 0$, and invests at time $t = 1$ as instructed at $t = 0$. The agent's degree of sophistication plays a key role in the delegation decision.

At $t = 0$, a naive agent plans for $t = 1$ to invest in stocks an amount

$$\phi_1^{\text{self, plan}} = \arg\max_{\phi} \mathbb{E}\big[u\big(w + (\tilde{x} - p)\,\phi\big)\big]. \tag{21}$$

If instead the agent were to delegate the investment decision, she would advise the manager to buy

$$\phi_1^{\text{delegate}} = \arg\max_{\phi} \mathbb{E}\big[u\big(w + (\tilde{x} - p)\,\phi - f\big)\big]. \tag{22}$$

Note that the agent evaluates the risk to occur at time $t = 1$ according to u, whether investment is delegated or not. When considering delegation at $t = 0$, the naive agent thus compares

$$\underbrace{\mathbb{E}\big[u\big(w + (\tilde{x} - p)\,\phi_1^{\text{delegate}} - f\big)\big]}_{\text{Naive with delegation}} \quad \text{vs.} \quad \underbrace{\mathbb{E}\big[u\big(w + (\tilde{x} - p)\,\phi_1^{\text{self, plan}}\big)\big]}_{\text{Naive without delegation}}. \tag{23}$$

The next result immediately follows.

Lemma 1. *Given a management fee $f > 0$, a naive HDRA agent never delegates the portfolio decision.*

The naive agent's comparison (23) is flawed, however. Once period $t = 1$ arrives, the risk is imminent and is evaluated according to the more risk averse v. Contrary to her plans at $t = 0$, the naive agent, if left to her own devices at $t = 1$, will only invest

$$\phi_1^{\text{self, actual}} = \arg\max_{\phi} \mathbb{E}\big[v\big(w + (\tilde{x} - p)\,\phi\big)\big]. \tag{24}$$

Since v is more risk averse than u, we know that $\phi_1^{\text{self, actual}} < \phi_1^{\text{self, plan}}$ (see Wang and Werner, 1994).

A sophisticated agent takes the future self's optimization problem as given and therefore optimizes subject to constraint (24). She thus compares

$$\underbrace{\mathbb{E}\big[u\big(w + (\tilde{x} - p)\,\phi_1^{\text{delegate}} - f\big)\big]}_{\text{Sophisticated with delegation}} \quad \text{vs.} \quad \underbrace{\mathbb{E}\big[u\big(w + (\tilde{x} - p)\,\phi_1^{\text{self, actual}}\big)\big]}_{\text{Sophisticated without delegation}}. \tag{25}$$

The left hand sides of the comparisons in (23) and (25) are the same; a naive and a sophisticated agent both correctly anticipate that a money manager will implement $\phi_1 = \phi_1^{\text{delegate}}$. However, the right hand sides of the comparisons differ since $\phi_1^{\text{self, actual}} < \phi_1^{\text{self, plan}}$. These comparisons lead to the following proposition.

Proposition 3. *Given HDRA preferences (v, u), a sophisticated agent delegates the portfolio decision if the management fee f is sufficiently small. Correspondingly, given a management fee f, a sophisticated agent with HDRA preferences (v, u) delegates the portfolio decision if v is sufficiently more risk averse than u.*

Consistent with the hypothesis that households choose financial advice to achieve a riskier portfolio allocation, Foerster et al. (2014) show that the foremost effect of financial advisors is to increase individuals' risky asset market participation (by up to 67%), as well as the extent of such risky asset market participation (by up to 39%). In addition, individuals who choose financial advice are less likely to close their accounts after negative return episodes. Finally, a comparison between these authors' OLS and IV estimates provides explicit support for the notion of demand for a commitment device to take risk: individuals who would otherwise not participate in the stock market are more likely to choose financial advice, a decision which subsequently increases their risky asset market participation.[6]

[6]To address the same puzzle, Gennaioli, Shleifer, and Vishny (2015) assume that agents delegate to "money doctors" because it reduces the perceived risk. Our model of anxiety predicts similar behavior based on a nonstandard preference rather than a belief distortion. Of course, effort costs of managing one's portfolio could also lead to delegation of investment management. However, effort costs cannot justify hiring a manager who underperforms the index on average, as buying index funds is virtually costless and free of effort. While buying the index is free of effort, but it is not free of short-term risk and associated "anxiety". Self 0 could thus correctly anticipate that the anxious self 1 will underperform the market even more than a random portfolio manager by failing to invest in equity at all. Self 0 will therefore be willing to pay an investment manager, even if she expects her to underperform the market.

An important assumption in deriving proposition 3 is that the agent cannot undo the delegation decision of period 0 once period 1 arrives. The commitment device that an anxiety-prone agent uses for risk-taking must be illiquid to some degree, similar to commitment devices a present-biased agent uses for saving, e.g. the "golden eggs" of Laibson (1997). Such illiquidity can be explicit, as in the case of Vanguard's "Personal Advisor Services" which go beyond just the behavioral coaching alluded to above. The service directly prevents clients from trading on their own (Vanguard Advisers, 2014). However, we also observe more subtle institutional features that provide illiquidity in arrangements in which risk-taking is delegated.

Fees are one obvious feature that discourages agents from undoing delegation arrangements once they are set up. HDRA thus provides one explanation for why redemption fees continue to feature prominently in the mutual fund industry, while management fees are increasingly being competed away (Khorana, Servaes, and Tufano, 2009). Different fees that also help agents commit to risk taking include, e.g. fees brokerage house charge for changing the ratio of equities and bonds in one's managed investment portfolio and fees that are imposed if the total exposure to a certain asset class falls below a threshold. The cost of having to provide liquidity without delay cannot necessarily explain such restrictions since the fund could easily charge the investor directly for such liquidation costs.

Another way to provide the desired illiquidity is to introduce delays. Especially high-risk forms of delegation such as hedge funds commonly impose initial lock-in periods and subsequent mandatory delays for withdrawals. Putting a temporal distance between the investor's decision to pull out and the valuation and payout of the investment prevents HDRA investors from "chickening out," as the delayed risky outcome is treated with lower risk aversion by the investor.

6 Term structure of risk prices

This section provides a stylized application of HDRA preferences to asset pricing. While the static setup of this paper has important limitations for this application, we want to illustrate the potential for HDRA to also account for asset pricing phenomena. We offer a dynamic model of HDRA specialized to asset pricing in Andries, Eisenbach, and Schmalz (2014).

A key feature of asset prices that has recently aroused much attention is a downward-sloping term structure of risk premia in the equity market (van Binsbergen et al., 2012, 2013). The authors empirically price a claim on the dividends of the S&P 500 index in the near future in contrast to the price of the S&P 500 itself which is a claim on *all* its future dividends. The risk-adjusted returns from holding the claim to only the short-term dividends are much higher than those from holding the claim to all future dividends. Specifically, van Binsbergen et al. (2012) show that the monthly Sharpe ratio of short-term dividend strips is 0.1124, almost twice as high as the monthly Sharpe ratio of the S&P 500 itself at 0.586. These results reflect that the price of near-term risk is significantly higher than the price of risk in the distant future. This finding raises a puzzle, because leading asset pricing models predict a nondecreasing term structure of risk premia.[7] We now show in a stylized financial market setting that HDRA can potentially account for this phenomenon.

We consider a setup in discrete time with two periods $t = 0, 1$ and two assets. Asset 0 pays a random dividend $\tilde{x}_0$ at the end of period 0 while asset 1 pays a random dividend $\tilde{x}_1$ at the end of period 1. Each asset is in net supply of 1 and the dividends $\tilde{x}_t$ are i.i.d. At the beginning of period 0, the agent has to form a portfolio (ϕ_0, ϕ_1) of the two assets as well as borrowing/lending a quantity ξ_t, for $t = 0, 1$. Given initial wealth w, the agent solves the following problem:

$$\max_{\{\phi_0,\phi_1,\xi_0,\xi_1\}} \mathbb{E}\big[v(\tilde{c}_0) + \delta u(\tilde{c}_1)\big] \tag{26}$$

$$\text{s.t.} \quad \tilde{c}_t = \tilde{x}_t\phi_t + \xi_t \quad \text{for} \quad t = 0, 1$$

$$p_0\phi_0 + \xi_0 + p_1\phi_1 + \frac{\xi_1}{1+r} \leq w.$$

Note that we have trading only at the beginning of period 0 to ensure that asset 1 is truly a long-term claim that is resolved and pays off in period 1. If we allowed for retrading at the beginning of period 1, asset 1 would become, first, a short-term claim on its price at the beginning of period 1 which, then, turns into a short-term claim on its dividend resolved and paid at the end of period 1. Because the single time of trading effectively provides commitment, naive and sophisticated behavior is

[7] Van Binsbergen et al. (2012) show that the term structure of risk premia is upward-sloping in the habit model of Campbell and Cochrane (1999) and the long-run risk model of Bansal and Yaron (2004) while it is flat in the rare disaster model of Gabaix (2012).

observationally equivalent in this section since both types of agents maximize the same objective function (26).

The first-order conditions for an interior solution to the problem (26) yield:

$$\begin{aligned} &\mathbb{E}\big[v'(\tilde{c}_0)\left(\tilde{x}_0 - p_0\right)\big] = 0 \\ \text{and} \quad &\mathbb{E}\big[\delta u'(\tilde{c}_1)\left(\tilde{x}_1 - (1+r)\,p_1\right)\big] = 0. \end{aligned} \tag{27}$$

For a mass of identical HDRA agents we have $\tilde{c}_0 = \tilde{x}_0$ and $\tilde{c}_1 = \tilde{x}_1$ which gives us the following result on risk premia.

Proposition 4. *If v is more risk averse than u, the risk premium on the short-term claim is higher than the risk premium on the long-term claim:*

$$\mathbb{E}[\tilde{x}_0] - p_0 > \mathbb{E}[\tilde{x}_1] - (1+r)\,p_1. \tag{28}$$

This result shows that in a very stylized setting, the HDRA model can account for the downward-sloping term structure of risk premia.

To understand the link between the Proposition 4 and the empirical evidence, it is important to bear in mind the distinction between a preference for early resolution of uncertainty and HDRA that we discuss in detail in Section 7.1. In the real world, immediate risk resolves for both short-term assets and long-term assets. In HDRA, however, not the horizon of resolution alone, but the horizon of the payouts affects utility. Only long-term assets, by definition, have risks that are associated with payouts in the distant future. Therefore, a preference for the early resolution of uncertainty cannot explain a downward-sloping term structure of risk premia.

The HDRA model also suggests an explanation for the value premium: as the duration of value stocks is shorter than that of growth stocks (Dechow et al., 2004) and HDRA investors dislike short-term risk, a value premium arises. Notably, this prediction arises directly from the utility specification and foregoes assumptions about the correlation structure of consumption growth and the stochastic discount factor (Lettau and Wachter, 2007).

While HDRA in this stylized setting suggests a potential solution to the puzzle of downward-sloping risk premia, the analysis has clear limitations. The static nature of the setup in this paper is not well suited for the inherently dynamic problem of asset pricing. Since we effectively assume a single time of trading at the beginning of period 0, we abstract from retrading which imposes additional no-arbitrage restrictions. A full analysis of these issues requires a setup that is more specialized toward asset pricing and less suited to model the other applications of HDRA illustrated in this paper. We therefore leave the asset pricing treatment to Andries, Eisenbach, and Schmalz (2014) where a generalization of non-time-separable Epstein and Zin (1989) preferences allows for a fully dynamic model with a separation of risk aversion and intertemporal substitution at different horizons. Hence both our static and our dynamic approach avoid confounding horizon-dependence in risk aversion in intertemporal substitution, an issue inherent in time-separable models of more than two periods (Section 3).

7 Comparison to and distinction from related theories

In this section we distinguish 'anxiety' from existing theories that are related orthogonal in concept or different in modeling approach and focus.

7.1 Preference for the timing of resolution of uncertainty

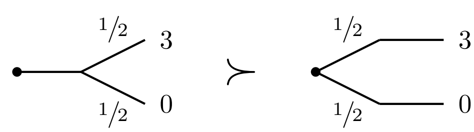

Figure 2: Preference for later resolution of uncertainty

The seminal paper by Kreps and Porteus (1978) is the first to consider a preference ranking between lotteries that differ in the timing of resolution of a given risk while the timing of the payoff is held constant. Fig. 2 illustrates such a preference. In contrast, HDRA manifests itself in comparisons of lotteries that are resolved and paid out at the same time and ranks them differently depending on temporal distance. Further, Kreps and Porteus (1978) explicitly rule out dynamically inconsistent behavior. In contrast, we allow for dynamic inconsistency.

Figure 3: Time-changing risk aversion

7.2 Time-changing risk aversion

A large literature in asset pricing assumes that agents' effective risk aversion changes over time, for example as a result of habit formation (Constantinides, 1990; Campbell and Cochrane, 1999).[8] Fig. 3 illustrates the choices of an agent who is more risk averse in one period than in another. In contrast, HDRA preferences are not time-varying. An anxious agent's effective risk aversion changes as a function of temporal distance to risk, not as a function of calendar time. Further, models of time-changing risk aversion are typically dynamically consistent.

7.3 Dynamically inconsistent time preferences

Agents with dynamic-inconsistency problems have been studied at least since Strotz (1955). Work in the tradition of Phelps and Pollak (1968) and Laibson (1997) focuses on inconsistent *time* preferences, e.g. modeled as quasi-hyperbolic discounting. The agent resolves *inter*temporal consumption tradeoffs differently depending on the time horizon: if the time horizon is short, the agent is more impatient than if the time horizon is long. We study an orthogonal dimension by assuming that the agent's *risk* preferences are dynamically inconsistent. The agent resolves *intra*temporal risk tradeoffs differently depending on the time horizon: if the horizon is short, the agent is more risk averse than if the horizon is long. Thus we emphasize that agents can be dynamically inconsistent independently in the dimensions of intertemporal consumption and intratemporal risk.

7.4 Temptation and self-control preferences

Motivated by the evidence on dynamically inconsistent time preferences, Gul and Pesendorfer (2001, 2004) propose an axiomatic model of temptation and self control. In the tradition of Kreps (1979), the model uses as primitive a preference relation over

[8]See Guiso et al. (2014) for an empirical study of time-changing risk aversion. Dillenberger and Rozen (2015) provide a model of history-dependent risk aversion.

choice sets and identifies an agent subject to temptation as one who prefers commitment, i.e. a subset of a choice set to the original set (see also Dekel et al., 2001). Under certain axioms, the preference relation over choice sets can be represented by a utility function consisting of a 'commitment utility' and a 'temptation utility.'

The representation suggests an agent who compromises between the commitment and the temptation utility when choosing *from* a choice set – possibly exerting costly self control – and who anticipates the cost of self control when choosing *among* choice sets.[9] The evidence on HDRA illustrated in Fig. 1 can be interpreted as an agent tempted by the safe alternative when choosing from the choice set {risky, safe} but preferring commitment to the risky alternative when choosing among the choice sets {risky}, {safe}.

However, the Gul-Pesendorfer agent is dynamically consistent by construction and precludes the analysis of agents with different levels of sophistication. For the case of dynamically inconsistent time preferences, O'Donoghue and Rabin (1999) generate rich implications by distinguishing different levels of sophistication. As illustrated by the parachute jump example in the introduction and our analysis of insurance markets, less than full sophistication about HDRA is not implausible. In general, we believe that agents can differ not only in their degree of HDRA but also in their level of sophistication about it. Our analysis therefore distinguishes between naive and sophisticated behavior, which we show has different implications for markets whose products are targeted towards different parts of the population. This approach is widely accepted in the literature on dynamically inconsistent time preferences (e.g. DellaVigna and Malmendier, 2006); we use it to analyze markets targeting consumers' dynamically inconsistent risk preferences. Appendix B shows that for situations in which sophistication does not matter, our approach generates the same results as the one by Gul-Pesendorfer.

7.5 Other theories

HDRA belongs with a set of theories that emphasize the impact of salience on decision making – temporal distance is but one dimension of salience. For example, in Bordalo, Gennaioli, and Shleifer (2012, 2013), the choice context makes certain aspects of lotteries more or less salient. This approach can account for several empirically rel-

[9]For the static model, Gul and Pesendorfer (2001) provide an axiomatization of an extended preference that justifies this interpretation.

evant phenomena that are different from those accounted for by HDRA. Epstein and Kopylov (2007) have a model of 'cold feet' in which agents become more *pessimistic* as risks approach, i.e. their subjective beliefs change. In contrast, HDRA is motivated by experimental evidence in which the objective probabilities are known to the subjects. Therefore, we keep beliefs fixed, but allow risk preferences to vary with the horizon. Epstein (2008) provides an axiomatization for a two-period model similar to ours. In contrast to our work, he uses the term 'anxiety' when an agent is more risk-averse for *distant* risks, evoking a notion of anticipatory feelings. Such anticipatory feelings are also an important aspect in Caplin and Leahy (2001) who expand the prize space to mental states and explain a set of economic phenomena different from the ones addressed in this paper. By contrast to this and other nonstandard preferences, the domain of our utility specification are standard intertemporal lotteries over monetary outcomes.

8 Conclusion

We model agents with horizon-dependent risk aversion who are more risk averse when risks are closer in time. The HDRA preference formulation we propose describes behavior of experimental subjects that cannot be described with existing modeling approaches and has rich implications across a wide range of applications in financial economics. In this paper, we show that HDRA can account for puzzles in consumer demand for insurance and in market prices of risk. We also show that sophistication about the dynamic risk inconsistency inherent in HDRA preferences and the associated utility costs can trigger institutional responses such as delegated portfolio management and illiquidity features of financial contracts.

While the present model is static to maintain transparency, the concept of HDRA is applicable more generally. In related work, Andries, Eisenbach, and Schmalz (2014) derive implications for general equilibrium asset pricing with horizon-dependent risk aversion in a fully dynamic model. Andries et al. (2015) document a key prediction of that model, a horizon-dependent price of variance risk. HDRA preferences also have implications for information acquisition and belief formation; Eisenbach and Schmalz (2015) show why it can be beneficial to a sophisticated anxiety-prone agent to hold overly precise beliefs and how such self-delusion can lead to excessive risk taking.

The present model's testable predictions include that more anxiety-prone individuals (those with greater disagreement between short-term and long-term risk aversion) exhibit a greater propensity for preference reversals in the face of approaching risks. It also predicts that the more sophisticated among these individuals exhibit a greater demand for commitment devices to take risk. Testing this prediction seems a promising road for future experimental and field research. Moreover, the model can be extended to generate testable implications also in other domains such as corporate finance, individual investor behavior, and household finance. We leave these extensions to future research.

Appendix

A Proofs

Proof of Proposition 1. Making use of the first-order condition (15) for I_0^*, the condition for buying additional insurance (17) can be rewritten as:

$$\mu_1 < \frac{\mu_0}{\frac{p_0-p_1}{1-p_1}\mu_0 + \frac{1-p_0}{1-p_1}} \tag{29}$$

Since $\mu_0 \geq 1$, the denominator on the right-hand side of (29) is weakly greater than 1 so we have

$$\mu_1 \leq \left(\frac{p_0 - p_1}{1 - p_1}\mu_0 + \frac{1 - p_0}{1 - p_1}\right)\mu_1 \tag{30}$$

We therefore have that condition (29) implies $\mu_1 < \mu_0$. Conversely, $\mu_1 > \mu_0$ implies that condition (29) is violated and the agent should not buy additional insurance. □

Proof of Proposition 2. From the proof of Proposition 1 we have that $\mu_1 > \mu_0$ implies

$$\frac{1}{\frac{p_0-p_1}{1-p_1} + \frac{1-p_0}{1-p_1}\frac{u'\left(A-\pi_0 I_0^*\right)}{u'\left(A-\pi_0 I_0^*-L+I_0^*\right)}} < \mu_1. \tag{31}$$

However, since v is more risk averse than u we have (by the same argument as in the proof of Proposition 4)

$$\frac{v'(A - \pi_0 I_0^* - L + I_0^*)}{v'(A - \pi_0 I_0^*)} > \frac{u'(A - \pi_0 I_0^* - L + I_0^*)}{u'(A - \pi_0 I_0^*)}, \tag{32}$$

and therefore

$$\frac{1}{\frac{p_0-p_1}{1-p_1} + \frac{1-p_0}{1-p_1}\frac{u'\left(A-\pi_0 I_0^*\right)}{u'\left(A-\pi_0 I_0^*-L+I_0^*\right)}} < \frac{1}{\frac{p_0-p_1}{1-p_1} + \frac{1-p_0}{1-p_1}\frac{v'\left(A-\pi_0 I_0^*\right)}{v'\left(A-\pi_0 I_0^*-L+I_0^*\right)}}, \tag{33}$$

with the difference increasing in the risk aversion of v relative to u. If v is sufficiently more risk averse than u, we therefore have

$$\frac{1}{\frac{p_0-p_1}{1-p_1}+\frac{1-p_0}{1-p_1}\frac{u'\left(A-\pi_0 I_0^*\right)}{u'\left(A-\pi_0 I_0^*-L+I_0^*\right)}}<\mu_1<\frac{1}{\frac{p_0-p_1}{1-p_1}+\frac{1-p_0}{1-p_1}\frac{v'\left(A-\pi_0 I_0^*\right)}{v'\left(A-\pi_0 I_0^*-L+I_0^*\right)}}, \tag{34}$$

i.e. the HDRA agent buys insurance although the non-HDRA agent does not. □

Proof of Lemma 1. The fund management fee f effectively reduces wealth and the agent is always worse off with lower wealth:

$$\frac{d}{df}\max_{\phi}\mathbb{E}\big[u\big(w+(\tilde{x}-p)\,\phi-f\big)\big]<0 \tag{35}$$

Given the definitions of ϕ_1^{delegate} and $\phi_1^{\text{self, plan}}$, the behavior of a naive agent follows immediately from (23) which is equivalent to the following inequality:

$$\max_{\phi}\mathbb{E}\big[u\big(w+(\tilde{x}-p)\,\phi-f\big)\big]<\max_{\phi}\mathbb{E}\big[u\big(w+(\tilde{x}-p)\,\phi\big)\big]. \tag{36}$$

Thus, a naive agent will never delegate. □

Proof of Proposition 3. Turning to a sophisticated agent, given the definition for $\phi_1^{\text{self, actual}}$ we have:

$$\max_{\phi}\mathbb{E}\big[u\big(w+(\tilde{x}-p)\,\phi\big)\big]>\mathbb{E}\big[u\big(w+(\tilde{x}-p)\,\phi_1^{\text{self, actual}}\big)\big]. \tag{37}$$

From condition (35) follows that there exists an $\bar{f}>0$ such that:

$$\max_{\phi}\mathbb{E}\big[u\big(w+(\tilde{x}-p)\,\phi-\bar{f}\big)\big]=\mathbb{E}\big[u\big(w+(\tilde{x}-p)\,\phi_1^{\text{self, actual}}\big)\big]. \tag{38}$$

For any $f\in[0,\bar{f}\big)$ we therefore have

$$\mathbb{E}\big[u\big(w+(\tilde{x}-p)\,\phi_1^{\text{delegate}}-f\big)\big]>\mathbb{E}\big[u\big(w+(\tilde{x}-p)\,\phi_1^{\text{self, actual}}\big)\big] \tag{39}$$

so the sophisticated agent will choose delegation. Given the results of Wang and Werner (1994), the inequality (37) is stronger the greater the difference in risk aversion between

v and u. Therefore the critical value $\bar{f}$ defined in (38) is smaller the greater the difference in risk aversion between v and u. □

Proof of Proposition 4. Since v is more risk averse than u we have

$$\begin{aligned} & -\frac{v''(x)}{v'(x)} > -\frac{u''(x)}{u'(x)} \\ \Rightarrow \quad & -\frac{d}{dx}\log v'(x) > -\frac{d}{dx}\log u'(x)\,. \end{aligned} \tag{40}$$

Integrating both sides over some interval $[a, b]$ yields

$$\frac{v'(b)}{v'(a)} < \frac{u'(b)}{u'(a)} \tag{41}$$

and the reverse inequality for $b < a$. For general a, b we therefore have

$$\left(\frac{u'(a)}{u'(b)} - \frac{v'(a)}{v'(b)}\right)(a - b) > 0. \tag{42}$$

Taking expectations for random $\tilde{a}$ we get

$$\frac{\mathbb{E}\big[u'(\tilde{a})\,(\tilde{a} - b)\big]}{u'(b)} > \frac{\mathbb{E}\big[v'(\tilde{a})\,(\tilde{a} - b)\big]}{v'(b)}. \tag{43}$$

Substituting in the price p_0 for b and the dividend $\tilde{x}$ for $\tilde{a}$ the RHS is zero by the first order conditions and we get

$$\mathbb{E}\big[u'(\tilde{x})\,(\tilde{x} - p_0)\big] > 0 \tag{44}$$

Using the first order conditions, this implies that $p_0 < (1 + r)\,p_1$ and since $\mathbb{E}[\tilde{x}_0] = \mathbb{E}[\tilde{x}_1]$ we have

$$\mathbb{E}[\tilde{x}_0] - p_0 > \mathbb{E}[\tilde{x}_1] - (1 + r)\,p_1 \tag{45}$$

as desired. □

B HDRA in the Gul-Pesendorfer framework

This appendix compares the implications of HDRA when analyzed in the framework of Gul and Pesendorfer (2001, 2004) are the same as in our framework. The Gul-

Pesendorfer representation yields a utility over choice set B given by

$$\mathcal{W}(B) = \max_{\tilde{x}\in B} \mathbb{E}\big[\mathcal{U}(\tilde{x}) + \mathcal{V}(\tilde{x})\big] - \max_{\tilde{y}\in B} \mathbb{E}\big[\mathcal{V}(\tilde{y})\big], \tag{46}$$

where $\mathcal{U}$ is the commitment utility and $\mathcal{V}$ the temptation utility. The representation suggests that when choosing from the choice set B, the agent maximizes

$$\hat{U}(\tilde{x}) = \mathbb{E}\big[\mathcal{U}(\tilde{x}) + \mathcal{V}(\tilde{x})\big], \tag{47}$$

i.e. a compromise between the commitment and the temptation utility. We are using the notation with hats to distinguish from the corresponding functions in our framework.

In our two-period setting, the commitment and temptation utility disagree about their risk aversion for $t = 0$ while agreeing about their risk aversion for $t = 1$ as well as intertemporal substitution. Analogous to the Gul-Pesendorfer implementation of impatience by Krusell et al. (2002, 2010), we can therefore implement HDRA by specifying

$$\mathcal{U}(c_0, c_1) = \hat{u}(c_0) + \delta\hat{u}(c_1), \tag{48}$$

$$\mathcal{V}(c_0, c_1) = \hat{v}(c_0) + \delta\hat{u}(c_1), \tag{49}$$

where $\hat{v}$ is more risk averse than $\hat{u}$. In $t = 0$, the agent maximizes

$$\hat{U}(\tilde{c}_0, \tilde{c}_1) = \mathbb{E}\big[\hat{u}(c_0) + \hat{v}(c_0) + 2\delta\hat{u}(c_1)\big]. \tag{50}$$

Note that with $\hat{v}$ more risk averse than $\hat{u}$, we also have $\hat{u} + \hat{v}$ more risk averse than $2\hat{u}$. This implementation of HDRA in the Gul-Pesendorfer framework therefore yields the same objective function for choice in $t = 0$ as our u-v framework if we set

$$\hat{u}(c) = \frac{1}{2}u(c), \tag{51}$$

$$\hat{v}(c) = v(c) - \frac{1}{2}u(c). \tag{52}$$

This implies that all of our results for sophisticated agents could also be generated in the Gul-Pesendorfer framework but not our results for naive agents.

References

Abdellaoui, M., E. Diecidue, and A. Onculer (2011). Risk preferences at different time periods: An experimental investigation. *Management Science 57*(5), 975–987.

Andries, M., T. M. Eisenbach, and M. C. Schmalz (2014). Asset pricing with horizon-dependent risk aversion. Working Paper.

Andries, M., T. M. Eisenbach, M. C. Schmalz, and Y. Wang (2015). The term structure of the price of volatility risk. Working Paper.

Baird, R. and D. Benady (1996). Money trap guarantees. *Marketing Week 19*, 32.

Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance 59*(4), 1481–1509.

Baucells, M. and F. Heukamp (2010). Common ratio using delay. *Theory and Decision 68*(1-2), 149–158.

van Binsbergen, J., M. Brandt, and R. Koijen (2012). On the timing and pricing of dividends. *American Economic Review 102*(4), 1596–1618.

van Binsbergen, J., W. Hueskes, R. Koijen, and E. Vrugt (2013). Equity yields. *Journal of Financial Economics 110*(3), 503–519.

Bordalo, P., N. Gennaioli, and A. Shleifer (2012). Salience theory of choice under risk. *Quarterly Journal of Economics 127*(3), 1243–1285.

Bordalo, P., N. Gennaioli, and A. Shleifer (2013). Salience and asset prices. *American Economic Review 103*(3), 623–628.

Breznitz, S. (2011). A study of worrying. *British Journal of Social and Clinical Psychology 10*(3), 271–279.

Campbell, J. Y. (2006). Household finance. *Journal of Finance 61*(4), 1553–1604.

Campbell, J. Y. and J. H. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy 107*(2), 205–251.

Caplin, A. and J. Leahy (2001). Psychological expected utility theory and anticipatory feelings. *Quarterly Journal of Economics 116*(1), 55–79.

Coble, K. and J. Lusk (2010). At the nexus of risk and time preferences: An experimental investigation. *Journal of Risk and Uncertainty 41*, 67–79.

Constantinides, G. M. (1990). Habit formation: A resolution of the equity premium puzzle. *Journal of Political Economy 98*(3), 519–543.

Dechow, P. M., R. G. Sloan, and M. T. Soliman (2004). Implied equity duration: A new measure of equity risk. *Review of Accounting Studies 9*(2-3), 197–228.

Dekel, E., B. L. Lipman, and A. Rustichini (2001). Representing preferences with a unique subjective state space. *Econometrica 69*(4), 891–934.

DellaVigna, S. and U. Malmendier (2006). Paying not to go to the gym. *American Economic Review 96*(3), 694–719.

Desai, P. S. and P. Padmanabhan (2004). Durable good, extended warranty and channel coordination. *Review of Marketing Science 2*(1).

Dew-Becker, I., S. Giglio, A. Le, and M. Rodriguez (2014). The price of variance risk. Working Paper.

Dillenberger, D. and K. Rozen (2015). History-dependent risk attitude. *Journal of Economic Theory 157*, 445–477.

Eisenbach, T. M. and M. C. Schmalz (2015). Anxiety, overconfidence, and excessive risk taking. Working Paper.

Eisner, R. and R. H. Strotz (1961). Flight insurance and the theory of choice. *Journal of Political Economy 69*(4), 355–368.

Epstein, L. G. (2008). Living with risk. *Review of Economic Studies 75*(4), 1121–1141.

Epstein, L. G. and I. Kopylov (2007). Cold feet. *Theoretical Economics 2*, 231–259.

Epstein, L. G. and S. E. Zin (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica 57*(4), 937–969.

Fenz, W. D. and S. Epstein (1967). Gradients of physiological arousal in parachutists as a function of an approaching jump. *Psychosomatic Medicine 29*(1), 33–51.

Fenz, W. D. and G. B. Jones (1972). Individual differences in physiologic arousal and performance in sport parachutists. *Psychosomatic Medicine 34*(1), 1–8.

Foerster, S., J. Linnainmaa, B. Melzer, and A. Previtero (2014). The costs and benefits of financial advice. Working Paper.

Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics 127*(2), 645–700.

Gennaioli, N., A. Shleifer, and R. W. Vishny (2015). Money doctors. *Journal of Finance 70*(1), 91–114.

Giglio, S., M. Maggiori, and J. Stroebel (2013). Very long-run discount rates. Working Paper.

Gollier, C. (2012). *Pricing the planet's future: The economics of discounting in an uncertain world*. Princeton University Press.

Gross, J. J. and L. F. Barrett (2011). Emotion generation and emotion regulation: One or two depends on your point of view. *Emotion Review 3*(1), 8–16.

Gruber, M. J. (1996). Another puzzle: The growth in actively managed mutual funds. *Journal of Finance 51*(3), 783–810.

Guiso, L., P. Sapienza, and L. Zingales (2014). Time varying risk aversion. Working Paper.

Gul, F. and W. Pesendorfer (2001). Temptation and self-control. *Econometrica 69*(6), 1403–1435.

Gul, F. and W. Pesendorfer (2004). Self-control and the theory of consumption. *Econometrica 72*(1), 119–158.

Holt, C. A. and S. K. Laury (2002). Risk aversion and incentive effects. *American Economic Review 92*(5), 1644–1655.

Jones, E. E. and C. A. Johnson (1973). Delay of consequences and the riskiness of decisions. *Journal of Personality 41*(4), 613–637.

Khorana, A., H. Servaes, and P. Tufano (2009). Mutual fund fees around the world. *Review of Financial Studies 22*(3), 1279–1310.

Kinniry, F. M., C. M. Jaconetti, M. A. DiJoseph, and Y. Zilbering (2014). *Putting a Value on Your Value: Quantifying Vanguard Advisor's Alpha*. The Vanguard Group.

Krantz, D. H., H. C. Kunreuther, et al. (2007). Goals and plans in decision making. *Judgment and Decision Making 2*(3), 137–168.

Kreps, D. M. (1979). A representation theorem for "preference for flexibility". *Econometrica 47*(3), 565–577.

Kreps, D. M. and E. L. Porteus (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica 46*(1), 185–200.

Krusell, P., B. Kurusçu, and A. A. Smith (2002). Time orientation and asset prices. *Journal of Monetary Economics 49*(1), 107–135.

Krusell, P., B. Kurusçu, and A. A. Smith, Jr. (2010). Temptation and taxation. *Econometrica 78*(6), 2063–2084.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics 112*(2), 443–477.

Laibson, D. (2015). Why don't present-biased agents make commitments? *American Economic Review 105*(5), 267–272.

Lettau, M. and J. A. Wachter (2007). Why is long-horizon equity less risky? A duration-based explanation of the value premium. *Journal of Finance 62*(1), 55–92.

Lo, A. W. and D. V. Repin (2002). The psychophysiology of real-time financial risk processing. *Journal of Cognitive Neuroscience 14*(3), 323–339.

Lo, A. W., D. V. Repin, and B. N. Steenbarger (2005). Fear and greed in financial markets: A clinical study of day-traders. *American Economic Review 95*(2), 352–359.

Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *Economic Journal 97*(387), 666–684.

Loewenstein, G. (1996). Out of control: Visceral influences on behavior. *Organizational Behavior and Human Decision Processes 65*(3), 272–292.

Loewenstein, G., E. U. Weber, C. K. Hsee, and E. Welch (2001). Risk as feelings. *Psychological Bulletin 127*(2), 267–286.

Luttmer, E. G. J. and T. Mariotti (2003). Subjective discounting in an exchange economy. *Journal of Political Economy 111*(5), 959–989.

MacKinlay, E. H. (1963). Airplane trip insurance. *Washington and Lee Law Review 20*(2), 346–353.

Maronick, T. J. (2007). Consumer perceptions of extended warranties. *Journal of Retailing and Consumer Services 14*(3), 224–231.

McClure, S., D. Laibson, G. Loewenstein, and J. Cohen (2004). Separate neural systems value immediate and delayed monetary rewards. *Science 306*(5695), 503–507.

Monat, A. (1976). Temporal uncertainty, anticipation time, and cognitive coping under threat. *Journal of Human Stress 2*(2), 32–43.

Monat, A. and R. S. Lazarus (1991). *Stress and Coping: An Anthology*. Columbia University Press.

Noussair, C. and P. Wu (2006). Risk tolerance in the present and the future: An experimental study. *Managerial and Decision Economics 27*(6), 401–412.

O'Donoghue, T. and M. Rabin (1999). Doing it now or later. *American Economic Review 89*(1), 103–124.

Onculer, A. (2000). Intertemporal choice under uncertainty: A behavioral perspective. Working Paper, INSEAD.

Paterson, R. J. and R. W. J. Neufeld (1987). Clear danger: Situational determinants of the appraisal of threat. *Psychological Bulletin 101*(3), 404–416.

Phelps, E. S. and R. A. Pollak (1968). On second-best national saving and game-equilibrium growth. *Review of Economic Studies 35*(2), 185–199.

Roth, W. T., G. Breivik, P. E. Jørgensen, and S. Hofmann (1996). Activation in novice and expert parachutists while jumping. *Psychophysiology 33*(1), 63–72.

Sagristano, M. D., Y. Trope, and N. Liberman (2002). Time-dependent gambling: Odds now, money later. *Journal of Experimental Psychology: General 131*(3), 364–376.

Schelling, T. C. (1984). Self-command in practice, in policy, and in a theory of rational choice. *American Economic Review 74*(2), 1–11.

Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies 23*(3), 165–180.

Trope, Y. and N. Liberman (2003). Temporal construal. *Psychological Review 110*(3), 403–421.

Vanguard Advisers, I. (2014). Vanguard personal advisor services brochure. http://www.adviserinfo.sec.gov/Iapd/Content/Common/crd_iapd_Brochure.aspx?BRCHR_VRSN_ID=253125.

Wang, Z. and J. Werner (1994). Portfolio characterization of risk aversion. *Economics Letters 45*(2), 259–265.

Weber, S. M. (2013). *Most Vanguard IRA Investors Shot Par by Staying the Course: 2008–2012*. The Vanguard Group.

Welch, E. (1999). The heat of the moment. Working Paper.