

History-Dependent Risk Attitude*

David Dillenberger[†]

University of Pennsylvania

Kareen Rozen[‡]

Yale University

March 2014

Abstract

We propose a model of history-dependent risk attitude, allowing a decision maker's risk attitude to be affected by his history of disappointments and elations. The decision maker recursively evaluates compound risks, classifying realizations as disappointing or elating using a threshold rule. We establish equivalence between the model and two cognitive biases: risk attitudes are reinforced by experiences (one is more risk averse after disappointment than after elation) and there is a primacy effect (early outcomes have the greatest impact on risk attitude). In a dynamic asset pricing problem, the model yields volatile, path-dependent prices.

Keywords: history-dependent risk attitude, reinforcement effect, primacy effect, dynamic reference dependence

JEL Codes: D03, D81, D91

*First version June 2010. This paper generalizes a previous version that circulated under the title "Disappointment Cycles." We benefitted from comments and suggestions by Simone Cerreia-Vioglio, Wolfgang Pesendorfer, Ben Polak, Andrew Postlewaite, Larry Samuelson, and several seminar audiences. We thank Xiaosheng Mu for excellent research assistance.

[†]Department of Economics, 160 McNeil Building, 3718 Locust Walk, Philadelphia, Pennsylvania 19104-6297. E-mail: ddill@sas.upenn.edu

[‡]Department of Economics and the Cowles Foundation for Research in Economics, 30 Hillhouse Avenue, New Haven, Connecticut 06511. E-mail: kareen.rozen@yale.edu. I thank the NSF for generous financial support through grant SES-0919955, and the economics departments of Columbia and NYU for their hospitality.

1 Introduction

Theories of decision making under risk typically assume that risk preferences are stable. Evidence suggests, however, that risk preferences may vary with personal experiences. It has been shown that emotions, which may be caused by exogenous factors or by the outcomes of past choices, play a large role in the decision to bear risk. Moreover, individuals are affected by unrealized outcomes, a phenomenon known in the psychological literature as counterfactual thinking.¹

Empirical work has found evidence of history-dependent risk aversion in a variety of fields. Pointing to adverse consequences for investment and the possibility of poverty traps, development economists have observed a long-lasting increase in risk aversion after natural disasters (Cameron and Shah, 2010) and, studying the dynamics of farming decisions in an experimental setting, increases of risk aversion after failures (Yesuf and Bluffstone, 2009). Malmendier and Nagel (2011) study how personal experiences of macroeconomic shocks affect financial risk-taking. Controlling for wealth, income, age, and year effects, they find that for up to three decades later, “households with higher experienced stock market returns express a higher willingness to take financial risk, participate more in the stock market, and conditional on participating, invest more of their liquid assets in stocks.” Applied work also demonstrates that changing risk aversion helps explain several economic phenomena. Barberis, Huang and Santos (2001) allow risk aversion to decrease with prior stock market gains (and increase with losses), and show that their model is consistent with the well-documented equity premium and excess volatility puzzles. Gordon and St-Amour (2000) study bull and bear markets, allowing risk attitudes to vary stochastically by introducing a state-dependent CRRA parameter in a discounted utility model. They show that countercyclical risk aversion best explains the cyclical nature of equity prices, suggesting that “future work should address the issue of determining the factors that underline the movements in risk preferences” which they identified.

In this work, we propose a model under which such shifts in risk preferences may arise. Our model of history-dependent risk attitude (HDRA) allows the way that risk unfolds over time to affect attitude towards further risk. We derive predictions for the comparative statics of risk aversion. In particular, our model predicts that one becomes more risk averse after a negative experience than after a positive one, and that sequencing matters: the earlier one is disappointed, the more risk averse one becomes.

¹On the effect of emotions, see Knutson and Greer (2008) and Kuhnen and Knutson (2011), as well as Section 1.1. Roesse and Olson (1995) offers a comprehensive overview of the counterfactual thinking literature.

To ease exposition, we begin by describing our HDRA model in the simple setting of T -stage, compound lotteries (the model is later extended to stochastic decision problems, in which the DM may take intermediate actions). A realization of a compound lottery is another compound lottery, which is one stage shorter. The DM categorizes each realization of a compound lottery as an elating or disappointing outcome. At each stage, his history is the preceding sequence of elations and disappointments. Each history h corresponds to a preference relation over one-stage lotteries. We consider one-stage preferences that are rankable in terms of their risk aversion. For example, an admissible collection could be a class of expected utility preferences with the Bernoulli function $u(x) = \frac{x^{1-\rho_h}}{1-\rho_h}$, where the coefficient of relative risk aversion ρ_h is history dependent.

The key features of the HDRA model are that compound lotteries are evaluated recursively and the DM's history assignment is *internally consistent*. More formally, starting at the final stage of the compound lottery and proceeding backwards, each one-stage lottery is replaced with its appropriate, history-dependent certainty equivalent. At each step of this recursive process, the DM is evaluating only one-stage lotteries, the outcomes of which are certainty equivalents of continuation lotteries. To determine which outcomes are elating and disappointing, the DM uses a threshold rule that assigns a number (a *threshold level*) to each one-stage lottery encountered in the recursive process. Internal consistency requires that if a subplottery is considered an *elating (disappointing)* outcome of its parent subplottery, then its certainty equivalent should exceed (or fall below) the threshold level corresponding to its parent subplottery. Mathematically, internal consistency imposes a fixed point requirement on the assignment of histories in a multi-stage setting.

Our model is general, allowing a wide class of preferences to be used for recursively evaluating lotteries, as well as a variety of threshold rules. The one-stage preferences may come from the betweenness class (Dekel, 1986; Chew, 1989), which includes expected utility as a special case. The DM's threshold rule may be either endogenous (preference-based) or exogenous. In the preference-based case, the DM's threshold moves endogenously with his preference; he compares the certainty equivalent of a subplottery to the certainty equivalent of its parent. In the exogenous case, the DM uses a rule that is independent of preferences but is a function of the lottery at hand; for example, an expectation-based rule that compares the certainty equivalent of a subplottery to his expected certainty equivalent. All of the components of the HDRA model – that is, the single-stage preferences, threshold rule, and history assignment – can be elicited from choice behavior.

Besides imposing internal consistency, we do not place any restriction on *how* risk aversion should depend on the history. Nonetheless, we show that the HDRA model predicts two well-documented cognitive biases; and that these biases are sufficient conditions for an HDRA representation to exist. First, the DM's risk attitudes are *reinforced* by prior experiences: he becomes

less risk averse after positive experiences and more risk averse after negative ones. Second, the DM displays a *primacy effect*: his risk attitudes are disproportionately affected by early realizations. In particular, the earlier the DM is disappointed, the more risk averse he becomes. We discuss evidence for these predictions in Section 1.1 below.

We show that the model and our characterization result readily extend to settings that allow for intermediate actions. As an application, we study a multi-period asset pricing problem where an HDRA decision maker (with CARA preferences) adjusts his asset holdings in each period after observing past dividend realizations. We show that the model yields volatile, path-dependent prices. Past realizations of dividends affect subsequent prices, even though they are statistically independent of future dividends and there are no income effects. For example, high dividends bring about price increases, while a sequence of only low dividends leads to an equity premium higher than in the standard, history-independent CARA case. Since risk aversion is endogenously affected by dividend realizations, the risk from holding an asset is magnified by expected future variation in the level of risk aversion. Hence the HDRA model introduces a channel of risk that is reflected in the greater volatility of asset prices. This is consistent with the observation of excess volatility in equity prices, dating to Shiller (1981).

This paper is organized as follows. Section 1.1 surveys evidence for the reinforcement and primacy effects. Section 1.2 discusses related literature. Section 2 introduces our model in the setting of compound lotteries (for notational simplicity, we extend the setting and results to stochastic decision trees only in Section 6). Section 3 contains our main result, which characterizes how risk aversion evolves with elations and disappointments. Section 4 describes how the components of the model can be elicited from choice behavior. Section 5 discusses further implications of the model. Section 6 generalizes the choice domain in the model to stochastic decision trees and studies a three-period asset pricing problem. Section 7 concludes. All proofs appear in the appendix.

1.1 Evidence for the Reinforcement and Primacy Effects

Our main predictions, the reinforcement and primacy effects, are consistent with a body of evidence on risk-taking behavior. Thaler and Johnson (1990) find that individuals become more risk averse after negative experiences and less risk averse after positive ones. Among contestants in the game show “Deal or No Deal,” Post, van den Assem, Baltussen and Thaler (2008) find mixed evidence, suggesting that contestants are more willing to take risks after extreme realizations. Guiso, Sapienza and Zingales (2011) estimate a marked increase in risk aversion in a sample of Italian investors after the 2008 financial crisis; the certainty equivalent of a risky gamble drops from 4,000 euros to 2,500, an increase in risk aversion which, as the authors show, cannot be due to changes

in wealth, consumption habits, or background risk. In an experiment with financial professionals, Cohn, Engelmann, Fehr and Maréchal (2013) find that subjects primed with a financial bust are significantly more risk averse than subjects primed with a boom. As discussed earlier, Malmendier and Nagel (2011) find that macroeconomic shocks lead to a long-lasting increase of risk aversion. Studying initial public offerings (IPOs), Kaustia and Knupfer (2008) identify pairs of “hot and cold” IPOs with close offer dates and follow the future subscription activities of investors whose first IPO subscription was in one of those two. They find that “twice as many investors participate in a subsequent offering if they first experience a hot offering rather than a cold offering.” Pointing to a primacy effect, they find that the initial outcome has a strong impact on subsequent offerings, and that “by the tenth offering, 65% of investors in the hot IPO group will have subscribed to another IPO, compared to only 39% in the cold IPO group.” Baird and Zelin (2000) study the impact of sequencing of positive and negative news in a company president’s letter. They find a primacy effect, showing that information provided early in the letter has the strongest impact on evaluations of that company’s performance. In general, sequencing biases² such as the primacy effect are robust and long-standing experimental phenomena (early literature includes Anderson (1965)); and several empirical studies, including Guiso, Sapienza and Zingales (2004) and Alesina and Fuchs-Schündeln (2007)), argue that early experiences may shape financial or cultural attitudes.

The biological basis of changes in risk aversion has been studied by neuroscientists. As summarized in Knutson and Greer (2008) and Kuhnen and Knutson (2011), neuroimaging studies have shown that two parts of the brain, the nucleus accumbens and the anterior insula, play a large role in risky decisions. The nucleus accumbens processes information on rewards, and is associated with positive emotions and excitement; while the anterior insula processes information about losses, and is associated with negative emotions and anxiety. Controlling for wealth and information, activation of the nucleus accumbens (anterior insula) is associated with bearing greater (lesser) risk in investment decisions. Discussing feedback effects, Kuhnen and Knutson (2011) note that:

[...] activation in the nucleus accumbens increases when we learn that the outcome of a past choice was better than expected (Delgado et al. (2000), Pessiglione et al. (2006)). Activation in the anterior insula increases when the outcome is worse than expected (Seymour et al (2004), Pessiglione et al (2006)), and when actions not chosen have larger payoffs than the chosen one.

In a neuroimaging study with 90 sequential investment decisions by subjects, these feedback effects are shown to influence subsequent risk-taking behavior (Kuhnen and Knutson, 2011).

²Another well-known sequencing bias is the recency effect, according to which more recent experiences have the strongest effect. A recency effect on risk attitude is opposite to the prediction of our model.

1.2 Relations to the literature

In many theories of choice over temporal lotteries, risk aversion can depend on the passage of time, wealth effects or habit formation in consumption; see Kreps and Porteus (1978), Segal (1990), Campbell and Cochrane (1999) and Rozen (2010), among others. We study how risk attitudes are affected by the past, independently of such effects. In the HDRA model, risk attitudes depend on “what might have been.” Such counterfactual thinking means that our model relaxes consequentialism (Machina, 1989; Hanany and Klibanoff, 2007), an assumption that is maintained by the papers above. Our form of history-dependence is conceptually distinct from models where current and future beliefs affect current utility (that is, dependence of utility on “what might be” in the future). This literature includes, for instance, Caplin and Leahy (2001) and Köszegi and Rabin (2009).

Caplin and Leahy (2001) propose a two-period model where the prize space of a lottery is enriched to contain psychological states, and there is an (unspecified) mapping from physical lotteries to mental states. Depending on how the second-period mental state is specified to depend on the first, Caplin and Leahy’s model could explain various types of risk-taking behaviors in the first period. While discussing the possibility of second-period disappointment, they do not address the question of history-dependence in choices. We conjecture that with additional periods and an appropriate specification of the mapping between mental states, one could replicate the predictions of our model. Köszegi and Rabin (2009) propose a utility function over T -period risky consumption streams. In their model, period utility is the sum of current consumption utility and the expectation of a gain-loss utility function, over all percentiles, of consumption utility at that percentile under the ex-post belief minus consumption utility at that percentile under the ex-ante belief. Beliefs are determined by an equilibrium notion, leading to multiplicity of possible beliefs. This bears resemblance to the multiplicity of internally consistent history assignments in our model (see Section 5 on how different assignments correspond to different attitudes to compound risks). Köszegi and Rabin (2009) do not address the question of history dependence: given an ex-ante belief over consumption, utility is not affected by prior history (how that belief was formed). While they point out that it would be realistic for comparisons to past beliefs to matter beyond one lag, they suggest one way to potentially model Thaler and Johnson (1990)’s result in their framework: “by assuming that a person receives money, and *in the same period* makes decisions on how to spend the money – with her old expectations still determining current preferences” (Köszegi and Rabin, 2009, Footnote 6). We conjecture that with additional historical differences in beliefs and an appropriate choice of functional forms (and relaxing additivity), one could replicate our predictions.

2 Framework

In this section we describe the essential components of our model of history dependent risk attitude (HDRA). Section 2.1 describes the domain of T -stage lotteries. Section 2.2 introduces the notion of history assignments. Section 2.3 discusses the recursive evaluation of compound lotteries. Section 2.4 introduces the key requirement of internal consistency, and formally defines the HDRA model.

2.1 Multi-stage lotteries: definitions and notations

Consider an interval of prizes $[w, b] \subset \mathbb{R}$. The choice domain is the set of T -stage simple lotteries over $[w, b]$. For any set X , let $\mathcal{L}(X)$ be the set of simple (i.e., finite support) lotteries over X . The set $\mathcal{L}^1 = \mathcal{L}([w, b])$ is the set of one-stage simple lotteries over $[w, b]$. The set $\mathcal{L}^2 = \mathcal{L}(\mathcal{L}^1)$ is the set of two-stage simple lotteries – that is, simple lotteries whose outcomes are themselves one-stage lotteries. Continuing in this manner, the set of T -stage simple lotteries is $\mathcal{L}^T = \mathcal{L}(\mathcal{L}^{T-1})$. A T -stage lottery could capture, for instance, an investment that resolves gradually, or a collection of monetary risks from different sources that resolve at different points in time.

We indicate the length of a lottery by its superscript, writing p^t , q^t , or r^t for an element of \mathcal{L}^t . A typical element p^t of \mathcal{L}^t has the form $p^t = \langle \alpha_1, p_1^{t-1}; \dots; \alpha_m, p_m^{t-1} \rangle$, which means that each $(t-1)$ -stage lottery p_j^{t-1} occurs with probability α_j . This notation presumes the outcomes are all distinct, and includes only those with $\alpha_j > 0$. For brevity, we sometimes write only $p^t = \langle \alpha_i, p_i^{t-1} \rangle_i$ for a generic t -stage lottery. One-stage lotteries are denoted by p, q and r , or simply $\langle \alpha_i, x_i \rangle_i$. At times, we use $p(x)$ to describe the probability of a prize x under the one-stage lottery p . For any $x \in X$, δ_x^ℓ denotes the ℓ -stage lottery yielding the prize x after ℓ riskless stages. Similarly, for any $p^t \in \mathcal{L}^t$, $\delta_{p^t}^\ell$ denotes the $(t+\ell)$ -stage lottery yielding the t -stage lottery p^t after ℓ riskless stages.

For any $t < \hat{t}$, we say that the t -stage lottery p^t is a *sublottery* of the \hat{t} -stage lottery $p^{\hat{t}}$ if there is a sequence $(p^\ell)_{\ell=t+1}^{\hat{t}-1}$ such that p^ℓ is in the support of $p^{\ell+1}$ for each $\ell \in \{t, \dots, \hat{t}-1\}$. In the case $\hat{t} = t+1$, this simply means that p^t is in the support of $p^{\hat{t}}$. By convention, we consider p^T a sublottery of itself. For any $p^T \in \mathcal{L}^T$, we let $S(p^T)$ be the set of all of its sublotteries.³

2.2 History assignments

Given a T -stage lottery p^T , the DM classifies each possible resolution of risk that leads from one sublottery to another as elating or disappointing. The DM's initial history is empty, denoted 0.

³Note that the same t -stage lottery p^t could appear in the support of different $(t+1)$ -stage sublotteries of p^T ; keeping this possibility in mind, but in order to economize on notation, throughout the paper we implicitly identify a particular sublottery p^t by the sequence of sublotteries leading to it. Viewed as a correspondence, we then have $S: \mathcal{L}^T \rightarrow \mathcal{P}(\bigcup_{t=1}^T \times_{\ell=T}^t \mathcal{L}^\ell)$, where \mathcal{P} is the power set.

If a sublottery p^t is degenerate – i.e., it leads to a given sublottery p^{t-1} with probability one – then the DM is not exposed to risk at that stage and his history is unchanged. If a sublottery p^t is nondegenerate, then the DM may be further elated (e) or disappointed (d) by the possible realizations. For any sublottery of p^T , the DM’s *history assignment* is given by his preceding sequence of elations and disappointments. Formally, the set of all possible histories is given by

$$H = \{0\} \cup \bigcup_{t=1}^{T-1} \{e, d\}^t. \quad (1)$$

The DM’s history assignment is a collection $a = \{a(\cdot|p^T)\}_{p^T \in \mathcal{L}^T}$, where for each $p^T \in \mathcal{L}^T$, the function $a(\cdot|p^T) : S(p^T) \rightarrow H$ assigns a history $h \in H$ to each sublottery of p^T , with the following restrictions. First, $a(p^T|p^T) = 0$. Second, the history assignment is sequential, in the sense that if p^{t+1} is a sublottery having p^t in its support, then $a(p^t|p^T) \in \{a(p^{t+1}|p^T)\} \times \{e, d\}$ if p^{t+1} is nondegenerate, and $a(p^t|p^T) = a(p^{t+1}|p^T)$ when p^{t+1} yields p^t with probability one.

Throughout the text we write he^t (or hd^t) to denote the concatenated history whereby e (or d) occurs t times after the history h . More generally, given two histories h and h' , we denote by hh' the concatenated history whereby h' occurs after h . These notations, wherever they appear, implicitly assume that the length of the resulting history is at most T (that is, it is still in H).

Example 1. Figure 1(a) considers the case $T = 3$, showing an example of a three-stage lottery p^3 and a history assignment. In the first stage of $p^3 = \langle .5, p^2; .5, \delta_q \rangle$, there is an equal chance to get either $\delta_q = \langle 1, q \rangle$ or $p^2 = \langle .25, p_1; .5, p_2; .25, p_3 \rangle$. The history assignment shown says $a(\delta_q|p^3) = a(q|p^3) = d$, $a(p^2|p^3) = e$, $a(p_1|p^3) = a(p_2|p^3) = ed$ and $a(p_3|p^3) = ee$.

2.3 Recursive evaluation of multi-stage lotteries

The DM evaluates multi-stage lotteries recursively, using history-dependent utility functions. Before describing the recursive process, we need to discuss the set of utility functions over one-stage lotteries (that is, the set \mathcal{L}^1) that will be applied. Let $\mathcal{V} = \{V_h\}_{h \in H}$ be the DM’s collection of utility functions, where each $V_h : \mathcal{L}^1 \rightarrow \mathbb{R}$ can depend on the DM’s history. In this paper, we confine attention to utility functions V_h that are continuous, monotone with respect to first-order stochastic dominance, and satisfy the following betweenness property (Dekel (1986), Chew (1989)):

Definition 1 (Betweenness). *The function V_h satisfies betweenness if for all $p, q \in \mathcal{L}^1$ and $\alpha \in [0, 1]$, $V_h(p) = V_h(q)$ implies $V_h(p) = V_h(\alpha p + (1 - \alpha)q) = V_h(q)$.*




Figure 1: In panel (a), a three-stage lottery p^3 and a history assignment. The first step of recursive evaluation yields the two-stage lottery in (b). The next step yields the one-stage lottery in (c).

Betweenness is a weakened form of the vNM-independence axiom: it implies neutrality toward randomization among equally-good lotteries, which retains the linearity of indifference curves in expected utility theory, but relaxes the assumption that they are parallel. This allows for a broad class of one-stage preferences which includes, besides expected utility, Gul's (1991) model of disappointment aversion and Chew's (1983) weighted utility. For each V_h , continuity and monotonicity ensure that any $p \in \mathcal{L}^1$ has a well-defined certainty equivalent, denoted $CE_h(p)$. That is, $CE_h(p)$ uniquely satisfies $V_h(\delta_{CE_h(p)}) = V_h(p)$.

The DM recursively evaluates each $p^T \in \mathcal{L}^T$ as follows. He first replaces each terminal, one-stage sublottery p with its history-dependent certainty equivalent $CE_{a(p|p^T)}(p)$. Observe that each two-stage sublottery $p^2 = \langle \alpha_i, p_i \rangle_i$ of p^T then becomes a one-stage lottery over certainty equivalents, $\langle \alpha_i, CE_{a(p_i|p^T)}(p_i) \rangle_i$, whose certainty equivalent itself can be evaluated using $CE_{a(p^2|p^T)}(\cdot)$. We now formally define the *recursive certainty equivalent* of any t -stage sublottery p^t , which we denote by $RCE(p^t|a, \mathcal{V}, p^T)$ to indicate that it depends on the history assignment a and the collection \mathcal{V} . In the case $t = 1$, the recursive certainty equivalent of p is simply its standard (history-dependent) certainty equivalent, $CE_{a(p|p^T)}(p)$. For each $t = 2, \dots, T$ and sublottery $p^t = \langle \alpha_i, p_i^{t-1} \rangle_i$ of p^T , the recursive certainty equivalent $RCE(p^t|a, \mathcal{V}, p^T)$ is given by

$$RCE(p^t|a, \mathcal{V}, p^T) = CE_{a(p^t|p^T)} \left(\left\langle \alpha_i, RCE(p_i^{t-1}|a, \mathcal{V}, p^T) \right\rangle_i \right). \quad (2)$$

Observe that the final stage gives the recursive certainty equivalent of p^T itself.

Example 1 (continued). Figure 1(b-c) shows how the lottery p^3 from (a) is recursively evaluated using the given history assignment. As shown in (b), the DM first replaces p_1, p_2, p_3 and q with their recursive certainty equivalents (which are simply the history-dependent certainty equiv-

alents). This reduces p^2 to a one-stage lottery, $\tilde{p} = \langle .25, CE_{ed}(p_1); .5, CE_{ed}(p_2); .25, CE_{ee}(p_3) \rangle$; and reduces δ_q to $\delta_{CE_d(q)}$. Following Equation (2), the recursive certainty equivalent of p^2 is $CE_e(\tilde{p})$; and the recursive certainty equivalent of δ_q is $CE_d(\delta_{CE_d(q)}) = CE_d(q)$. These replace \tilde{p} and $\delta_{CE_d(q)}$, respectively, in (c). The resulting one-stage lottery is then evaluated using $CE_0(\cdot)$ to find the recursive certainty equivalent of p^3 .

2.4 A model of history-dependent risk attitude

Recall that at each step of the recursive process described above, the DM is evaluating only one-stage lotteries, the outcomes of which are recursive certainty equivalents of continuation lotteries. The HDRA model requires the history assignment to be *internally consistent* in each step of this process. Roughly speaking, for the DM to consider p_j^t , an elating (or disappointing) outcome of its parent lottery $p^{t+1} = \langle \alpha_i, p_i^t \rangle_i$, the recursive certainty equivalent of p_j^t , $RCE(p_j^t | a, \mathcal{V}, p^T)$, must fall above (below) a *threshold level* that depends on $\langle \alpha_i, RCE(p_i^t | a, \mathcal{V}, p^T) \rangle_i$. Note that in the recursive process, the threshold rule acts on folded-back lotteries (not ultimate prizes). To formalize this, let \bar{H} be the set of non-terminal histories in H , that is, taking the union only up to $T - 2$ in Equation (1). Allowing the threshold level to depend on the DM's history-dependent risk attitude, a threshold rule is a function $\tau : \bar{H} \times \mathcal{L}^1 \rightarrow [w, b]$ such that for each $h \in \bar{H}$, $\tau_h(\cdot)$ is continuous, monotone, satisfies betweenness (see Definition 1), and has the feature that for any $x \in [w, b]$, $\tau_h(\delta_x) = x$.

Definition 2 (Internal consistency). *The history assignment $a = \{a(\cdot | p^T)\}_{p^T \in \mathcal{L}^T}$ is internally consistent given the threshold rule τ and collection \mathcal{V} if for every p^T , and for any p_j^t in the support of a nondegenerate sublottery $p^{t+1} = \langle \alpha_i, p_i^t \rangle_i$, the history assignment $a(\cdot | p^T)$ satisfies the following property: if $a(p^{t+1} | p^T) = h \in \bar{H}$ and $a(p_j^t | p^T) = he$ (hd), then⁴*

$$RCE(p_j^t | a, \mathcal{V}, p^T) \geq (<) \tau_h \left(\langle \alpha_1, RCE(p_1^t | a, \mathcal{V}, p^T); \dots; \alpha_n, RCE(p_n^t | a, \mathcal{V}, p^T) \rangle \right).$$

We consider two types of threshold-generating rules different DMs may use: an *exogenous* rule (independent of preference) and an *endogenous*, preference-based specification:

Exogenous threshold. For some f from the betweenness class, the threshold rule τ is independent of h and implicitly given for any $p \in \mathcal{L}^1$ by $f(p) = f(\delta_{\tau_h(p)})$. For example, if f is an expected

⁴We assume a DM considers an outcome of a nondegenerate sublottery elating if its recursive certainty equivalent is at least as large as the threshold of the parent lottery. Alternatively, it would not affect our results if we instead assume an outcome is disappointing if its recursive certainty equivalent is at least as small as the threshold, or even introduce a third assignment, neutral (n), which treats the case of equality differently than elation or disappointment. In any case, a generic nonempty history consists of a sequence strict elations and disappointments.

utility functional for some increasing $u : [w, b] \rightarrow \mathbb{R}$, then $\tau(p) = u^{-1}(\sum u(x)p(x))$. If u is also linear, then $\tau(p) = \mathbb{E}(p)$. In this case we refer to τ as an expectation-based rule.

Note that an exogenous threshold rule $\tau_h(\cdot)$ is independent of the collection \mathcal{V} , even though it ultimately takes as an input lotteries that have been generated recursively using \mathcal{V} . For example, in the case of expectation-based rule the DM compares the certainty equivalent of a sublottery to his expected certainty equivalent.

Endogenous, preference-based threshold. The threshold rule τ is given by $\tau_h(\cdot) = CE_h(\cdot)$. In this case, the DM's history-dependent risk attitude affects his threshold for elation and disappointment, and the condition for internal consistency reduces to comparing the recursive certainty equivalent of p'_j with that of its parent sublottery p^{t+1} . (The reason for the name "endogenous" is that the threshold depends on the DM's current preference, which is itself determined endogenously.)

To illustrate the difference between the two types of threshold rules, consider a one-stage lottery giving prizes $0, 1, \dots, 1000$ with equal probabilities. If the DM is risk averse and uses the (endogenous) preference-based threshold rule, then he may be elated by prizes smaller than 500, where the cutoff for elation is his certainty equivalent for this lottery. By contrast, if he uses the exogenous threshold rule, then only prizes exceeding 500 are elating. That is, exogenous threshold rules separate the classification of disappointment and elation from preferences.

We now present the HDRA model, which determines a utility function $U(\cdot|a, \mathcal{V}, \tau)$ over \mathcal{L}^T using the recursive certainty equivalent from Equation (2).

Definition 3 (History-dependent risk attitude, HDRA). *An HDRA representation over T -stage lotteries consists of a collection $\mathcal{V} := \{V_h\}_{h \in H}$ of utilities over one-stage lotteries from the betweenness class, a history assignment a , and an (endogenous or exogenous) threshold rule τ , such that the history assignment a is internally consistent given τ and \mathcal{V} , and we have for any p^T ,*

$$U(p^T|a, \mathcal{V}, \tau) = RCE(p^T|a, \mathcal{V}, p^T).$$

We identify a DM with an HDRA representation by the triple (a, \mathcal{V}, τ) satisfying the above.

It is easy to see that the HDRA model is ordinal in nature: the induced preference over \mathcal{L}^T is invariant to increasing, potentially different transformations of the members of \mathcal{V} . This is because the HDRA model takes into account only the certainty equivalents of sublotteries after each history.

In the HDRA model, the DM's risk attitudes depend on the prior sequence of disappointments and elations, but not on the "intensity" of those experiences. That is, the DM is affected only by his general impressions of past experiences. This simplification of histories can be viewed as an extension of the notions of elation and disappointment for one-stage lotteries suggested in Gul (1991)

or Chew (1989). In those works, a prize x is an elating outcome of a lottery p if it is preferred to p itself, and is a disappointing outcome if p is preferred to it. We generalize the threshold rule for elation and disappointment, and apply this idea recursively throughout the compound lottery to classify sublotteries as elating or disappointing. While the classification of a realization is binary, the probabilities and magnitudes of realizations affect the value of the threshold for elation and disappointment, and in general affect the utility of the lottery. By permitting risk attitude to depend only on prior elations and disappointments, this specification allows us to study endogenously evolving risk attitudes under a parsimonious departure from history independence. Behaviorally, such a classification of histories may describe a cognitive limitation on the part of the DM. The DM may find it easier to recall whether he was disappointed or elated, than whether he was very disappointed or slightly elated. Keeping track of the “exact” intensity of disappointment and elation for every realization – which is itself a compound lottery – may be difficult, leading the DM to classify his impressions into discrete categories: sequences of elations and disappointments.⁵

To illustrate the model’s internal consistency requirement, we return to Example 1.

Example 1 (continued). *Suppose the DM uses an expectation-based exogenous threshold rule $\tau_h(\cdot) = \mathbb{E}(\cdot)$ and the assignment from Figure 1(a). To verify that it is internally consistent to have $a(p_1|p^3) = a(p_2|p^3) = ed$ and $a(p_3|p^3) = ee$, one must check that*

$$CE_{ee}(p_3) \geq \mathbb{E}\left(\langle .25, CE_{ed}(p_1); .5, CE_{ed}(p_2); .25, CE_{ee}(p_3) \rangle\right) > CE_{ee}(p_1), CE_{ee}(p_2),$$

since the recursive certainty equivalents of p_1, p_2, p_3 are given by the corresponding standard certainty equivalents. Next, recall that the recursive certainty equivalent of p^2 is given by $CE_e(\tilde{p})$, where $\tilde{p} = \langle .25, CE_{ed}(p_1); .5, CE_{ed}(p_2); .25, CE_{ee}(p_3) \rangle$. To verify that it is internally consistent to have $a(p^2|p^3) = e$ and $a(\delta_q|p^3) = d$, one must check that

$$CE_e(\tilde{p}) \geq \mathbb{E}\left(\langle .5, CE_e(\tilde{p}); .5, CE_d(q) \rangle\right) > CE_d(q).$$

Observe that internal consistency imposes a fixed point requirement that takes into account the entire history assignment for a lottery. In Example 1, even if p_1, p_2, p_3 have an internally consistent assignment within p^2 , this assignment must lead to a recursive certainty equivalent for p^2 that is

⁵A stylized assumption in the model is that the DM treats a period which is completely riskless differently than a period in which any amount of risk resolves. This implies, in particular, that receiving a lottery with probability one is treated discontinuously differently than receiving a “nearby” sublottery with elating and disappointing outcomes. This simplifying assumption may be descriptively plausible in situations as described above, where the DM only recalls whether he was disappointed, elated, or neither (since he was not exposed to any risk). Alternatively, this may relate to situations where emotions are triggered by the mere “possibility” of risk (see a discussion of the phenomenon of “probability neglect” in Sunstein (2002)).

elating relative to that of δ_q . We next explore the implications internal consistency has for risk attitudes.

3 Characterization of HDRA

In this section, we investigate for which collections of single-stage preferences \mathcal{V} and threshold rule τ can an HDRA representation (a, \mathcal{V}, τ) exist. In the case that the single-stage preferences are history independent (say, $V_h = V$ for all $h \in H$), an internally consistent assignment can always be constructed, because the recursive certainty equivalent of a sublottery is independent of the assigned history.⁶ To study what happens when risk attitudes are shaped by prior experience, and for the sharpest characterization of when an internally consistent history assignment exists in this case, we consider collections \mathcal{V} for which the utility functions after each history are (strictly) rankable in terms of their risk aversion.

Definition 4. We say that V_h is strictly more risk averse than $V_{h'}$, denoted $V_h >_{RA} V_{h'}$, if for any $x \in X$ and any nondegenerate $p \in \mathcal{L}^1$, $V_h(p) \geq V_h(\delta_x)$ implies that $V_{h'}(p) > V_{h'}(\delta_x)$.

Definition 5. We say that the collection $\mathcal{V} = \{V_h\}_{h \in H}$ is ranked in terms of risk aversion if for all $h, h' \in H$, either $V_h >_{RA} V_{h'}$ or $V_{h'} >_{RA} V_h$.

Examples of \mathcal{V} with the rankability property are a collection of expected CRRA utilities, $\mathcal{V} = \{\mathbb{E}(\frac{x^{1-\rho_h}}{1-\rho_h} | \cdot)\}_{h \in H}$, or a collection of expected CARA utilities, $\mathcal{V} = \{\mathbb{E}(1 - e^{-\rho_h x} | \cdot)\}_{h \in H}$, with distinct coefficients of risk aversion (i.e., $\rho_h \neq \rho_{h'}$ for all $h, h' \in H$). A non-expected utility example is a collection of Gul's (1991) disappointment aversion preferences, with history-dependent coefficients of disappointment aversion, β_h .⁷ In all these examples, history-dependent risk aversion is captured by a single parameter. Under the rankability condition (Definition 5), we now show that the existence of an HDRA representation implies regularity properties on \mathcal{V} that are related to

⁶The assignment can be constructed by labeling a sublottery elating (disappointing) whenever its fixed, history-independent recursive certainty equivalent is greater than (resp., weakly smaller than) the corresponding threshold.

⁷According to Gul's model, the value of a lottery p , $V(p; \beta, u)$ is the unique v solving

$$v = \frac{\sum_{\{x|u(x) \geq v\}} p(x)u(x) + (1 + \beta) \sum_{\{x|u(x) < v\}} p(x)u(x)}{1 + \beta \sum_{\{x|u(x) < v\}} p(x)}, \quad (3)$$

where $u : X \rightarrow \mathbb{R}$ is increasing and $\beta \in (-1, \infty)$ is the coefficient of disappointment aversion. According to (3), lotteries are evaluated by calculating their "expected utility," except that disappointing outcomes (those that are worse than the lottery) get a uniformly greater (or smaller) weight depending on β . Gul shows that the DM becomes increasingly risk averse as the disappointment aversion coefficient increases, holding the utility over prizes constant. An admissible collection is thus $\mathcal{V} = \{V(\cdot; \beta_h, u)\}_{h \in H}$, where $V(\cdot; \beta_h, u)$ is given by (3) and the coefficients β_h are distinct.

well-known cognitive biases; and that these properties imply the existence of an HDRA representation.

3.1 The reinforcement and primacy effects

Experimental evidence suggests that risk attitudes are reinforced by prior experiences. They become less risk averse after positive experiences and more risk averse after negative ones. This effect is captured in the following definition.

Definition 6. $\mathcal{V} = \{V_h\}_{h \in H}$ displays the reinforcement effect if $V_{hd} >_{RA} V_{he}$ for all h .

A body of evidence also suggests that individuals are affected by the position of items in a sequence. One well-documented cognitive bias is the *primacy effect*, according to which early observations have a strong effect on later judgments. In our setting, the order in which elations and disappointments occur affect the DM’s risk attitude. The primacy effect suggests that the shift in attitude from early realizations can have a lasting and disproportionate effect. Future elation or disappointment can mitigate, but not overpower, earlier impressions, as in the following definition.

Definition 7. $\mathcal{V} = \{V_h\}_{h \in H}$ displays the primacy effect if $V_{hdet} >_{RA} V_{hedt}$ for all h and t .

The reinforcement and primacy effects together imply strong restrictions on the collection \mathcal{V} , as seen in the following observation. We refer below to the lexicographic order on histories of the same length as the ordering where \tilde{h} precedes h if it precedes it alphabetically. Since d comes before e , this is interpreted as “the DM is disappointed earlier in \tilde{h} than in h .”

Observation 1. \mathcal{V} displays the reinforcement and primacy effects if and only if for h, \tilde{h} of the same length, $V_{\tilde{h}} >_{RA} V_h$ if \tilde{h} precedes h lexicographically. Moreover, under the additional assumption $V_{hd} >_{RA} V_h >_{RA} V_{he}$ for all $h \in H$, \mathcal{V} displays the reinforcement and primacy effects if and only if for any h, h', h'' , we have $V_{hdh''} >_{RA} V_{heh'}$.

The content of Observation 1 is visualized in Figure 2. The first statement corresponds to the lexicographic ordering within each row. Under the additional assumption $V_{hd} >_{RA} V_h >_{RA} V_{he}$, which says an elation reduces (and a disappointment increases) the DM’s risk aversion *relative to his initial level*, one obtains the vertical lines and consecutive row alignment. Observe that along a realized path, this imposes no restriction on how current risk aversion compares to risk aversion two or more periods ahead *when the continuation path* consists of both elating and disappointing outcomes: e.g., one can have either $V_h <_{RA} V_{hed}$ or $V_h >_{RA} V_{hed}$.

We are now ready to state the main result of the paper.

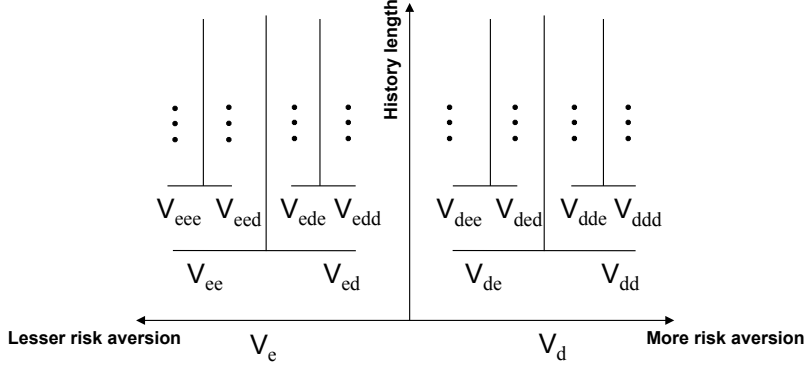


Figure 2: Starting from the bottom, each row depicts the risk aversion rankings $>_{RA}$ of the V_h for histories of length $t = 1, 2, 3, \dots, T - 1$. The reinforcement and primacy effects imply the lexicographic ordering in each row. The vertical boundaries and consecutive row alignment would be implied by the additional assumption $V_{hd} >_{RA} V_h >_{RA} V_{he}$ for all $h \in H$.

Theorem 1 (Necessary and sufficient conditions for HDRA). *Consider a collection \mathcal{V} that is ranked in terms of risk aversion, and an exogenous or endogenous threshold rule τ . An HDRA representation (a, \mathcal{V}, τ) exists – that is, there exists an internally consistent assignment a – if and only if \mathcal{V} displays the reinforcement and primacy effects.*

The proof appears in the Appendix. Observe that the model takes as given a collection of preferences \mathcal{V} that are ranked in terms of risk aversion, but does not specify *how* they are ranked. Theorem 1 shows that internal consistency places strong restrictions on how risk aversion evolves with elation and disappointment.

4 Eliciting the Components of the HDRA Model

In this section, we study how to recover the primitives (a, \mathcal{V}, τ) from the choice behavior of a DM who applies the HDRA model. Let \succeq denote the DM's observed preference ranking over \mathcal{L}^T .

First, the utility functions in \mathcal{V} may be elicited using only choice behavior over the simple subclass of lotteries illustrated in Figure 3(a) and defined as follows. Let

$$\mathcal{L}_u^T = \left\{ \langle \alpha_1, \langle \alpha_2, \dots \langle \alpha_{T-1}, p; 1 - \alpha_{T-1}, \delta_{z_{T-1}} \rangle \dots; 1 - \alpha_2, \delta_{z_2}^{T-2}; 1 - \alpha_1, \delta_{z_1}^{T-1} \rangle \right\}$$

such that $p \in \mathcal{L}^1$, $z_i \in \{b, w\}$, and $\alpha_i \in [0, 1]$

be the set of lotteries where in each period, either the DM learns he will receive one of the extreme prizes b or w for sure, or he must incur further risk (which is ultimately resolved by p if an extreme prize has not been received). For lotteries in the class \mathcal{L}_u^T the history assignment is unambiguous.




Figure 3: A representative lottery in (a) \mathcal{L}_u^T and (b) \mathcal{M}_u^T , with $p, q \in \mathcal{L}^1$ and $z_i \in \{b, w\}$ for all i .

The DM is disappointed (elated) by any continuation sublottery received *instead* of the best prize b (resp., the worst prize w). To illustrate how the class \mathcal{L}_u^T allows elicitation of \mathcal{V} , consider a history $h = (h_1, \dots, h_t)$ of length $t \leq T - 1$; as a convention, the length t of the initial history $h = 0$ is zero. Pick any sequence $\alpha_1, \dots, \alpha_t \in (0, 1)$ and take $\alpha_i = 1$ for $i > t$. (Note that anytime the continuation probability α_i is one, the history is unchanged). Construct the sequence of prizes z_1, \dots, z_t such that for every $i \leq t$, $z_i = b$ if $h_i = d$, and $z_i = w$ if $h_i = e$. Finally, define $\ell_h : \mathcal{L}^1 \rightarrow \mathcal{L}_u^T$ by $\ell_h(r) \equiv \langle \alpha_1, \langle \alpha_2, \dots \langle \alpha_t, \delta_r^{T-t-1}; 1 - \alpha_t, \delta_{z_t}^{T-t} \rangle \dots; 1 - \alpha_2, \delta_{z_2}^{T-2} \rangle; 1 - \alpha_1, \delta_{z_1}^{T-1} \rangle$ for any $r \in \mathcal{L}^1$, with the convention that $\delta_r^0 = r$. It is easy to see that the history assignment of r must be h . Moreover, under the HDRA model, $V_h(p) \geq V_h(q)$ if and only if $\ell_h(p) \succeq \ell_h(q)$. Note that only the ordinal rankings represented by the collection \mathcal{V} affect choice behavior.

We may also elicit the DM's (endogenous or exogenous) threshold rule τ from his choices. To do this, it again suffices to examine the DM's rankings over a special class of T -stage lotteries. Similarly to the class \mathcal{L}_u^T defined above, we can define the class \mathcal{M}_u^T of T -stage lotteries having the form illustrated in Figure 3(b): in the first $t \leq T - 2$ stages, the DM either receives an extreme prize or faces additional risk, which is finally resolved by a two stage lottery of the form

$$m(p, q) = \left\langle \frac{1}{2}, p; \frac{q(x_1)}{2}, \delta_{x_1}; \dots, \frac{q(x_n)}{2}, \delta_{x_n} \right\rangle, \quad (4)$$

where $p, q \in \mathcal{L}^1$, with the notation $q = \langle q(x_1), x_1; \dots; q(x_n), x_n \rangle$. Moreover, similarly to the construction of $\ell_h(\cdot)$, we may define for any history $h = (h_1, \dots, h_t) \in \bar{H}$ the function $m_h : \mathcal{L}^1 \times \mathcal{L}^1 \rightarrow \mathcal{M}_u^T$ as follows. For any nondegenerate $p, q \in \mathcal{L}^1$,

$$m_h(p, q) \equiv \langle \alpha_1, \langle \alpha_2, \dots \langle \alpha_t, \delta_{m(p,q)}^{T-t-2}; 1 - \alpha_t, \delta_{z_t}^{T-t} \rangle \dots; 1 - \alpha_2, \delta_{z_2}^{T-2} \rangle; 1 - \alpha_1, \delta_{z_1}^{T-1} \rangle,$$

where the continuation probabilities $(\alpha_i)_{i=1}^t \in (0, 1)$ and extreme prizes $(z_i)_{i=1}^t \in \{w, b\}$ are selected so that the history assignment of $m(p, q)$ is precisely h .

We can determine how a prize z compares to the threshold value of a one-stage lottery q as follows. Suppose for the moment that there exists a nondegenerate $p \in \mathcal{L}^1$ satisfying $CE_{he}(p) = z$ and for which $m_h(p, q) \sim m_h(\delta_z, q)$. We claim that the history assignment of p in $m_h(p, q)$ cannot be hd . Indeed, $CE_{hd}(p) < CE_{he}(p) = z$ implies that $m_h(p, q) \sim m_h(\delta_z, q) \succ m_h(\delta_{CE_{hd}(p)}, q)$, where the strict preference follows because a two-stage lottery of the form $m(\delta_x, q)$ is isomorphic to the one-stage lottery $\langle \frac{1}{2}, x; \frac{q(x_1)}{2}, x_1; \dots; \frac{q(x_n)}{2}, x_n \rangle$, to which monotonicity of V_h then applies. Therefore, the history assignment of p in $m_h(p, q)$ must be he , which means, by internal consistency, that

$$z = CE_{he}(p) \geq \tau_h(\langle \frac{1}{2}, CE_{he}(p); \frac{q(x_1)}{2}, x_1; \dots; \frac{q(x_n)}{2}, x_n \rangle). \quad (5)$$

The one-stage lottery $\langle \frac{1}{2}, z; \frac{q(x_1)}{2}, x_1; \dots; \frac{q(x_n)}{2}, x_n \rangle$ is a convex combination of the sure prize z and the lottery q . Since τ_h satisfies betweenness, Equation (5) holds if and only if $z \geq \tau_h(q)$. Similarly, if for some nondegenerate $p \in \mathcal{L}^1$ satisfying $CE_{hd}(p) = z$, $m_h(p, q) \sim m_h(\delta_z, q)$, then $z < \tau_h(q)$. So far we have only assumed that a p with the desired properties exists whenever $z \geq (<) \tau_h(q)$; we show in the Appendix that this is indeed the case. Thus, we can represent τ_h 's comparisons between a lottery q and a sure prize z through the auxiliary relation \succeq_{τ_h} , defined by $\delta_z \succeq_{\tau_h} q$ ($q \succ_{\tau_h} \delta_z$) if there is a nondegenerate $p \in \mathcal{L}^1$ such that $m_h(p, q) \sim m_h(\delta_z, q)$ and $CE_{he}(p) = z$ (resp., $CE_{hd}(p) = z$). Proposition 2 in the Appendix shows how to complete this relation using choice over lotteries in \mathcal{M}_u^T , and proves that it represents the threshold τ_h .

Finally, to recover the history assignment of a T -stage lottery, one needs to iteratively ask the DM what sure outcomes should replace the terminal lotteries to keep him indifferent. Generically, his chosen outcome must be the certainty equivalent of the corresponding subplottery under his history assignment a .⁸

5 Other Properties of HDRA

Optimism and pessimism. An HDRA representation is identified by the triple (a, \mathcal{V}, τ) . Given a threshold rule τ and a collection \mathcal{V} satisfying the reinforcement and primacy effects, Theorem 1 guarantees that an internally consistent history assignment exists. There may in fact be more than one internally consistent assignment for some lotteries, meaning that it is possible for two HDRA decision-makers to agree on \mathcal{V} and τ but to sometimes disagree on which outcomes are

⁸Since H is finite, if there is p^T such that two assignments yield the same value, then there is an open ball around p^T within which every other lottery has the property that no two assignments yield the same value.

elating and disappointing. For a simple example, consider $T = 2$ and suppose that $CE_e(p) > z > CE_d(p)$; then it is internally consistent for p to be called elating or disappointing in $\langle \alpha, p; 1 - \alpha, \delta_z \rangle$. Therefore a is not a redundant primitive of the model. We can think of some plausible rules for generating history assignments. For example, given the pair (\mathcal{V}, τ) , the DM may be an *optimist* (*pessimist*) if for each p^T he selects the internally consistent history assignment a that maximizes (minimizes) the HDRA utility of p^T . We can then say that if the collections \mathcal{V}^A and \mathcal{V}^B are ordinally equivalent and $\tau^A = \tau^B$, then $(a^A, \mathcal{V}^A, \tau^A)$ is *more optimistic* than $(a^B, \mathcal{V}^B, \tau^B)$ if $U(p^T | a^A, \mathcal{V}^A, \tau^A) \geq U(p^T | a^B, \mathcal{V}^B, \tau^B)$ for every p^T . Suppose we define a comparative measure of compound risk aversion by saying that DM^A is *less compound-risk averse* than DM^B if for any $p^T \in \mathcal{L}^T$ and $x \in [w, b]$, $U(p^T | a^B, \mathcal{V}^B, \tau^B) \geq U(\delta_x^T | a^B, \mathcal{V}^B, \tau^B)$ implies $U(p^T | a^A, \mathcal{V}^A, \tau^A) \geq U(\delta_x^T | a^A, \mathcal{V}^A, \tau^A)$. It is immediate to see that if one DM is more optimistic than another, then he is less compound-risk averse. As pointed out in Section 1.2, the multiplicity of possible internally consistent history assignments, and the use of an assignment rule to pick among them, resembles the multiplicity of possible beliefs in Köszegi and Rabin (2009), and their use of the “preferred personal equilibrium” criterion.

Statistically reversing risk attitudes. A second implication of Theorem 1, in the case of an endogenous threshold rule, is statistically reversing risk attitudes. (This feature does not arise under exogenous threshold rules.) When the threshold moves with preference, a DM who has been elated is not only less risk averse than if he had been disappointed, but also has a higher elation threshold. In other words, the reinforcement effect implies that after a disappointment, the DM is more risk averse and “settles for less”; whereas after an elation, the DM is less risk averse and “raises the bar.” Therefore, the probability of elation in any sublottery increases if that sublottery is disappointing instead of elating.⁹ Moreover, the “mood swings” of a DM with an endogenous threshold need not moderate with experience, even under the additional assumption shown in Figure 2 that $V_{hd} >_{RA} V_h >_{RA} V_{he}$ for all h . Indeed, suppose the DM’s risk aversion is described by a collection of risk aversion coefficients $\{\rho_h\}_{h \in H}$. For any fixed T , the parameters need not satisfy $|\rho_{ed} - \rho_e| \geq |\rho_{ede} - \rho_{ed}| \geq |\rho_{eded} - \rho_{ede}| \cdots$.

⁹The psychological literature, in particular Parducci (1995) and Smith, Diener and Wedell (1989), provides support for this prediction regarding elation thresholds. Summarizing these works, Schwarz and Strack (1998) observe that “an extreme negative (positive) event increased (decreased) satisfaction with subsequent modest events....Thus, the occasional experience of extreme negative events facilitates the enjoyment of the modest events that make up the bulk of our lives, whereas the occasional experience of extreme positive events reduces this enjoyment.”

Preferring to lose late rather than early: the “second serve” effect. A recent New York Times article¹⁰ documents a widespread phenomenon in professional tennis: to avoid a double fault after missing the first serve, many players employ a “more timid, perceptibly slower” second serve that is likely to get the ball in play but leaves them vulnerable in the subsequent rally. In the article, Daniel Kahneman attributes this to the fact that “people prefer losing late to losing early.” Kahneman says that “a game in which you have a 20 percent chance to get to the second stage and an 80 percent chance to win the prize at that stage...is less attractive than a game in which the percentages are reversed.” Such a preference was first noted in Ronen (1973).

To formalize this, take $\alpha \in (.5, 1)$ and any two prizes $H > L$. How does the two-stage lottery $p_{late}^2 = \langle \alpha, \langle 1 - \alpha, H; \alpha, L \rangle; 1 - \alpha, \delta_L \rangle$, where the DM has a good chance of delaying losing, compare with $p_{early}^2 = \langle 1 - \alpha, \langle \alpha, H; 1 - \alpha, L \rangle; \alpha, \delta_L \rangle$, where the DM is likely to lose earlier? (For simplicity we need only consider two stages here, but to embed this into \mathcal{L}^T we may consider each of p_{late}^2 and p_{early}^2 as a sublottery evaluated under a nonterminal history $h \in \bar{H}$, similarly to the construction in \mathcal{M}_u^T .) Standard expected utility predicts indifference over p_{late}^2 and p_{early}^2 , because the distribution over final outcomes is the same. To examine the predictions of the HDRA model for the case that each $V_h \in \mathcal{V}$ is from the expected utility class, let u_h denote the Bernoulli utility corresponding to utility V_h . In both lotteries, reaching the final stage is elating, since $H > L$. The HDRA value of p_{late}^2 is higher than the value of p_{early}^2 starting from history h if and only if

$$\alpha u_h \left(u_{he}^{-1} \left((1 - \alpha) u_{he}(H) + \alpha u_{he}(L) \right) \right) + (1 - \alpha) u_h(L) > (1 - \alpha) u_h \left(u_{he}^{-1} \left(\alpha u_{he}(H) + (1 - \alpha) u_{he}(L) \right) \right) + \alpha u_h(L). \quad (6)$$

Proposition 1. *The DM prefers losing late to losing early (that is, Equation (6) holds for all $H > L$, $\alpha \in (.5, 1)$ and $h \in \bar{H}$) if and only if $V_{he} <_{RA} V_h$.*

The proof appears in the Appendix, where we show that Equation (6) is equivalent to u_h being a concave transformation of u_{he} . Similarly, one also can show the equivalence between preferring to win sooner rather than later, and $V_{hd} >_{RA} V_h$.

Nonmonotonic behavior: thrill of winning and pain of losing. A DM with an HDRA representation may violate first-order stochastic dominance for certain compound lotteries. For example, again for the case $T = 2$, if α is very high, the lottery $\langle \alpha, p; 1 - \alpha, \delta_w \rangle$ may be preferred to $\langle \alpha, p; 1 - \alpha, \delta_b \rangle$; in the former, p is evaluated as an elation, while in the latter, it is evaluated as

¹⁰See “Benefit of Hitting Second Serve Like the First,” August 29, 2010, available for download at <http://www.nytimes.com/2010/08/30/sports/tennis/30serving.html?pagewanted=all>.

a disappointment. Because the prizes w and b are received with very low probability, the “thrill of winning” the lottery p may outweigh the “pain of losing” the lottery p . This arises from the reinforcement effect on compound risks. While monotonicity with respect to compound first-order stochastic dominance may be normatively appealing, the appeal of such monotonicity is rooted in the assumption of consequentialism (that “what might have been” does not matter). As Mark Machina points out, once consequentialism is relaxed, as is explicitly done in this paper, violations of monotonicity may naturally occur.¹¹ In our model, violations of monotonicity arise only on particular compound risks, in situations where the utility gain or loss from a change in risk attitude outweighs the benefit of a prize itself. The idea that winning is enjoyable and losing is painful may also translate to nonmonotonic behavior in more general settings. For example, Lee and Malmendier (2011) show that forty-two percent of auctions for board games end at a price which is higher than the simultaneously available buy-it-now price.

6 Extension to intermediate actions, and a dynamic asset pricing problem

In this section, we extend our previous results to settings where the DM may take intermediate actions while risk resolves. We then apply the model to a three-period asset pricing problem to examine the impact of history-dependent risk attitude on prices.

6.1 HDRA with intermediate actions

In the HDRA model with intermediate actions, the DM categorizes each realization of a dynamic (stochastic) decision problem – which is a choice set of shorter dynamic decision problems – as elating or disappointing. He then recursively evaluates all the alternatives in each choice problem based on the preceding sequence of elations and disappointments.

Formally, for any set Z , let $\mathcal{K}(Z)$ be the set of finite, nonempty subsets of Z . The set of one-stage decision problems is given by $\mathcal{D}^1 = \mathcal{K}(\mathcal{L}(X))$. By iteration, the set of t -stage decision problem is given by $\mathcal{D}^t = \mathcal{K}(\mathcal{L}(\mathcal{D}^{t-1}))$. A one-stage decision problem $D^1 \in \mathcal{D}^1$ is simply a set of one-stage lotteries. A t -stage decision problem $D^t \in \mathcal{D}^t$ is a set of elements of the form $\langle \alpha_i, D_i^{t-1} \rangle_i$, each of which is a lottery over $(t-1)$ -stage decision problems. Note that our earlier domain of T -stage lotteries can be thought of as the subset of \mathcal{D}^T where all choice sets are degenerate.

¹¹As discussed in Mas-Colell, Whinston and Green (1995), Machina offers the example of a DM who would rather take a trip to Venice than watch a movie about Venice, and would rather watch the movie than stay home. Due to the disappointment he would feel watching the movie in the event of not winning the trip itself, Machina points out that the DM might prefer a lottery over the trip and staying home, to a lottery over the trip and watching the movie.

The admissible collections of one-stage preferences $\mathcal{V} = \{V_h\}_{h \in H}$ and threshold rules are unchanged. The set of possible histories H is also the same as before, with the understanding that histories are now assigned to choice nodes. For each $D^T \in \mathcal{D}^T$, the history assignment $a(\cdot|D^T)$ maps each choice problem within D^T to a history in H that describes the preceding sequence of elations and disappointments. The initial history is empty, i.e. $a(D^T|D^T) = 0$,

The DM recursively evaluates each T -stage decision problem D^T as follows. For a terminal decision problem D^1 , the recursive certainty equivalent is simply given by $RCE(D^1|a, \mathcal{V}, x^T) = \max_{p \in D^1} CE_{a(D^1|D^T)}(p)$. That is, the value of the choice problem D^1 is the value of the ‘best’ lottery in it, calculated using the history corresponding to D^1 . For $t = 2, \dots, T$, the recursive certainty equivalent is

$$RCE(D^t|a, \mathcal{V}, D^T) = \max_{\langle \alpha_i, D_i^{t-1} \rangle_{i \in D^t}} CE_{a(D^t|D^T)}(\langle \alpha_i, RCE(D_i^{t-1}|a, \mathcal{V}, D^T) \rangle).$$

This is analogous to the definition of the recursive certainty equivalent from before, with the addition of the max operator that indicates that the DM chooses the best available continuation problem. The history assignment of choice sets must be internally consistent. Given a decision problem D^T , if D_j^t is an elating (disappointing) outcome of $\langle \alpha_1, D_1^t; \dots; \alpha_n, D_n^t \rangle \in D^{t+1}$, then it must be that $RCE(D_j^t|a, \mathcal{V}, D^T) \geq (<) \tau_{a(D^{t+1}|D^T)}(\langle \alpha_1, RCE(D_1^t|a, \mathcal{V}, D^T); \dots; \alpha_n, RCE(D_n^t|a, \mathcal{V}, D^T) \rangle)$. The definition of HDRA imposes the same requirements as before, but over this larger choice domain.

Definition 8 (HDRA with intermediate actions). *An HDRA representation over T -stage decision problems consists of a collection $\mathcal{V} := \{V_h\}_{h \in H}$ of utilities over one-stage lotteries from the betweenness class, a history assignment a , and an (endogenous or exogenous) threshold rule τ , such that the history assignment a is internally consistent given τ and \mathcal{V} , and we have for any D^T ,*

$$U(D^T|a, \mathcal{V}, \tau) = RCE(D^T|a, \mathcal{V}, D^T).$$

Observe that the DM is ‘‘sophisticated’’ under HDRA with intermediate actions. From any future choice set, the DM anticipates selecting the best continuation decision problem. That choice leads to an internally consistent history assignment of that choice set. When reaching a choice set, the single-stage utility he uses to evaluate the choices therein is the one he anticipated using, and his choice is precisely his anticipated choice. Internal consistency is thus a stronger requirement than before, because it takes optimal choices into account. However, our previous result extends.

Theorem 2 (Extension to intermediate actions). *Consider a collection \mathcal{V} that is ranked in terms of risk aversion, and an exogenous or endogenous threshold rule τ . An HDRA representation with*

intermediate actions (a, \mathcal{V}, τ) exists – that is, there exists an internally consistent assignment a – if and only if \mathcal{V} displays the reinforcement and primacy effects.

6.2 Application to a three-period asset pricing problem

We now apply the model to study a three-period asset-pricing problem in a representative agent economy. We confine attention to three periods because it is the minimal time horizon T needed to capture both the reinforcement and primacy effects. We show that the model yields predictable, path-dependent prices that exhibit excess volatility arising from actual and anticipated changes in risk aversion.

In each period $t = 1, 2, 3$, there are two assets traded, one safe and one risky. At the end of the period, the risky asset yields \tilde{y} , which is equally likely to be *High* (H) or *Low* (L). The second is a risk-free asset returning $R = 1 + r$, where r is the risk-free rate of return. Asset returns are in the form of a perishable consumption good that cannot be stored; it must be consumed in the current period. Each agent is endowed with one share of the risky asset in each period. The realization of the risky asset in period t is denoted y_t . In the beginning of each period $t > 1$, after a sequence of realizations (y_1, \dots, y_{t-1}) , each agent can trade in the market, at price $P(y_1, \dots, y_{t-1})$ for the risky asset, with the risk-free asset being the numeraire. At $t = 1$, there are no previous realizations and the price is simply denoted P . At the end of each period, the agent learns the realization of \tilde{y} and consumes the perishable return. The payoff at each terminal node of the three-stage decision problem is the sum of per-period consumptions.¹² In each period t , the agent's problem is to determine the share $\alpha(y_1, \dots, y_{t-1})$ of property rights to retain on his unit of risky asset given the asset's past realizations (y_1, \dots, y_{t-1}) . The agent is purchasing additional shares when $\alpha(y_1, \dots, y_{t-1}) > 1$, and is short-selling when $\alpha(y_1, \dots, y_{t-1}) < 0$. At $t = 1$, there are no previous realizations and his share is simply denoted α .

The agent has HDRA preferences with underlying CARA expected utilities; that is, the Bernoulli function after history h is $u_h(x) = 1 - e^{-\lambda_h x}$. Notice that none of the results in this application depend on whether the DM employs an exogenous or endogenous threshold rule (as there are only two possible realizations, *High* or *Low*, of the asset in each stage). The CARA specification of one-stage preferences means that our results will not arise from wealth effects. For this section, we use the following simple parametrization of the agent's coefficients of absolute risk aversion. Consider a, b satisfying $0 < a < 1 < b$ and $\lambda_0 > 0$. In the first period, elation scales down the agent's risk aversion by a^2 , while disappointment scales it up by b^2 . In the second period, elation

¹²Alternatively, one could let the terminal payoff be some function of the consumption vector, in which case the DM is evaluating lotteries over terminal utility instead of total consumption.

scales down the agent's current risk aversion by a , while disappointment scales it up by b . In summary, $\lambda_e = a^2\lambda_0$, $\lambda_d = b^2\lambda_0$, $\lambda_{ee} = a^3\lambda_0$, $\lambda_{ed} = a^2b\lambda_0$, $\lambda_{de} = b^2a\lambda_0$, and $\lambda_{dd} = b^3\lambda_0$. This parametrization satisfies the reinforcement and primacy effects, and has the feature that $\lambda_h \in (\lambda_{he}, \lambda_{hd})$ for all h . As can be seen from our analysis below, if the agent's risk aversion is independent of history and fixed at λ_0 at every stage, then the asset price is constant over time. By contrast, in the HDRA model, prices depend on past realizations of the asset, even though past and future realizations are statistically independent. The following result formalizes the predictions of the HDRA model.

Theorem 3. *Given the parametrization above, the price responds to past realizations as follows:*

- (i) $P(H,H) > P(H,L) > P(L,H) > P(L,L)$ at $t = 3$.
- (ii) $P(H) > P(L)$ at $t = 2$.
- (iii) Price increases after each *High* realization: $P(H,H) > P(H) > P$ and $P(L,H) > P(L)$.
- (iv) P , $P(L)$, and $P(L,L)$ are all below the (constant) price under history independent risk aversion λ_0 .

Theorem 3 is illustrated in Figure 4, which depicts the simulated price path for the specification $R = 1$, $\lambda_0 = .005$, $b = 1.2$, $a = .8$, $H = 20$, and $L = 0$.¹³ In that case, the price also decreases after each *Low* realization: $P > P(L) > P(L,L)$ and $P(H) > P(H,L)$.¹⁴ In general, this need not be true. Notice that the DM faces one fewer stage of risk each time there is a realization of the asset. Nonetheless, price is constant with history-independent CARA preferences. With history-dependent risk aversion, a compound risk may become even riskier due to the fact that the continuation certainty equivalents fluctuate with risk aversion. That is, expected future risk aversion movements introduce an additional source of risk, causing price volatility. The price after each history is a convex combination of H and L , with weights that depend on the product of current risk aversion and the spread between the future certainty equivalents (as seen in Table 1). Depending on how much risk aversion fluctuates, there may be an upward trend in prices, simply from having fewer stages of risk left. Elation (*High* realizations) reinforces that trend,

¹³Estimates of CARA coefficients in the literature are highly variable, ranging from .00088 (Cohen and Einav (2007)) to .0085 to .14 (see Saha, Shumway and Talpaz (1994) for a summary of estimates). Using $\lambda_0 = .005$, $b = 1.2$ and $a = .8$ yields CARA coefficients between .0025 and .00864.

¹⁴Barberis, Huang and Santos (2001) propose and calibrate a model where investors have linear loss aversion preferences and derive gain-loss utility only over fluctuations in financial wealth. In our model, introducing consumption shocks would induce shifts in risk aversion. They assume that the amount of loss aversion decreases with a statistic that depends on past stock prices. In their calibration, this leads to price increases (decreases) after good (bad) dividends and high volatility.




Figure 4: The predicted HDRA price path when $R = 1$, $\lambda_0 = .005$, $a = .8$, $b = 1.2$, $H = 20$, and $L = 0$. The price in the standard model is given by \bar{P} in Equation (7) below.

because the agent is both less risk averse and faces fewer stages of risk. However, there is tension between these two forces after disappointment (*Low* realizations), because the agent is more risk averse even though he faces a shorter horizon. One can find parameter values where the upward trend dominates, and *Low* realizations yield a (quantitatively very small) price increase – while maintaining the rankings in Theorem 3. Intuitively, this occurs when disappointment has a very weak effect on risk aversion ($b \approx 1$) but elation has a strong effect, because then expected variability in future risk aversion (hence expected variability in utility) prior to a realization may overwhelm the small increase in risk aversion after a *Low* realization occurs.

The proof of Theorem 3 is in the Appendix. There, we show how the agent uses the HDRA model to rebalance his portfolio. We first solve the agent’s optimization problem under the recursive application of one-stage preferences using an arbitrary history assignment. We later find that the only internally consistent assignment has the DM be elated each time the asset realization is *High*, and disappointed each time it is *Low*. Since in equilibrium the agent must hold his initial share of the asset, the conditions from portfolio optimization pin down prices given a history assignment.

To better understand the implications of HDRA for prices, it is useful to first think about the standard setting with history-independent risk aversion. If risk aversion were fixed at λ_0 , then the price \bar{P} in the standard model (the dotted line in Figure 4) would simply be constant and given by

$$\bar{P} = \frac{1}{R} \left(\frac{1}{1 + \exp(\lambda_0(H - L))} H + \frac{\exp(\lambda_0(H - L))}{1 + \exp(\lambda_0(H - L))} L \right), \quad (7)$$

as can be seen from Equations (9)-(11) in the Appendix. In the HDRA model, however, the asset price depends on y_1 and y_2 , not only through current risk aversion, but also through the impact on future certainty equivalents. Let $h(y_1, \dots, y_{t-1})$ be the agent's history assignment after the realizations y_1, \dots, y_{t-1} . Thus the agent's current risk aversion is given by $\lambda_{h(y_1, \dots, y_{t-1})}$. To describe how prices evolve with risk aversion, we introduce one additional piece of notation. Given a one-dimensional random variable \tilde{x} and a function ϕ of that random variable, we let $\Gamma_{\tilde{x}}(\lambda, \phi(\tilde{x}))$ denote the certainty equivalent of $\phi(\tilde{x})$ given CARA preferences with coefficient λ . That is,

$$\Gamma_{\tilde{x}}(\lambda, \phi(\tilde{x})) = -\frac{1}{\lambda} \ln \mathbb{E}_{\tilde{x}} \left[\exp(-\lambda \phi(\tilde{x})) \right].$$

Then, our analysis in the Appendix shows that the asset price in period t takes the form

$$P(y_1, \dots, y_{t-1}) = \frac{1}{R} \left(\frac{1}{1 + \exp(f(y_1, \dots, y_{t-1}))} H + \frac{\exp(f(y_1, \dots, y_{t-1}))}{1 + \exp(f(y_1, \dots, y_{t-1}))} L \right), \quad (8)$$

where the weighting function f is given by the following table:

Realizations	Weight $f(y_1, \dots, y_{t-1})$
0	$\lambda_0(H - L + \Gamma_{\tilde{y}_2}(\lambda_e, \tilde{y}_2 + \Gamma_{\tilde{y}_3}(\lambda_{h(H, \tilde{y}_2)}, \tilde{y}_3)) - \Gamma_{\tilde{y}_2}(\lambda_d, \tilde{y}_2 + \Gamma_{\tilde{y}_3}(\lambda_{h(L, \tilde{y}_2)}, \tilde{y}_3)))$
y_1	$\lambda_{h(y_1)}(H - L + \Gamma_{\tilde{y}_3}(\lambda_{h(y_1)e}, \tilde{y}_3) - \Gamma_{\tilde{y}_3}(\lambda_{h(y_1)d}, \tilde{y}_3))$
y_1, y_2	$\lambda_{h(y_1, y_2)}(H - L)$

Table 1: The weighting function $f(y_1, \dots, y_{t-1})$ for prices.

As can be seen from Equation (8), an increase in $f(y_1, \dots, y_{t-1})$ decreases the weight on H and thus decreases the price of the asset. All the price comparisons in Theorem 3 follow from Equation (8) and Table 1, with some comparisons easier to see than others. For instance, the ranking $P(H, H) > P(H, L) > P(L, H) > P(L, L)$ in Theorem 3(i) follows immediately from the reinforcement and primacy effects and the fact that $H > L$. The ranking $P(H) > P(L)$ in Theorem 3(ii) is proved in Lemma 6 in the Appendix. To see why some argument is required, notice that $\lambda_e < \lambda_d$ is not sufficient to show that $P(H) > P(L)$ unless we know something more about how the difference in the certainty equivalents of period-three consumption, $\Gamma_{\tilde{y}_3}(\lambda_{h(y_1)e}, \tilde{y}_3) - \Gamma_{\tilde{y}_3}(\lambda_{h(y_1)d}, \tilde{y}_3)$, compares after $y_1 = H$ versus $y_1 = L$. To see how the rankings $P(H, H) > P(H)$ and $P(L, H) > P(L)$ in Theorem 3(iii) follow from Table 1 above, two observations are needed. First, the parameterization implies $\lambda_{ee} < \lambda_e$, and $\lambda_{de} < \lambda_d$. Moreover, the term $\Gamma_{\tilde{y}_3}(\lambda_{h(y_1)e}, \tilde{y}_3) - \Gamma_{\tilde{y}_3}(\lambda_{h(y_1)d}, \tilde{y}_3)$ is positive, as will be shown in our argument for internal consistency. An additional step of proof, given in

Lemma 7 of the Appendix, is needed to show that $P(H) > P$. Similarly, Theorem 3(iv) follows from $\lambda_{dd} > \lambda_d > \lambda_0$ combined with the fact that λ_h always multiplies a term strictly larger than $H - L$. Hence the prices P , $P(L)$, and $P(L,L)$ all fall below the price \bar{P} in the standard model with constant risk aversion λ_0 .

7 Concluding remarks

We propose a model of history-dependent risk attitude which has tight predictions for how disappointments and elations affect the attitude to future risks. The model permits a wide class of preferences and threshold rules, and is consistent with a body of evidence on risk-taking behavior. To study endogenous reference dependence under a minimal departure from recursive history independent preferences, HDRA posits the categorization of each sublottery as either elating or disappointing. The DM's risk attitudes depend on the prior sequence of disappointments or elations, but not on the "intensity" of those experiences.

It is possible to generalize our model so that the more a DM is "surprised" by an outcome, the more his risk aversion shifts away from a baseline level. The equivalence between the generalized model and the reinforcement and primacy effects remains.¹⁵ Extending the model requires introducing an additional component (a *sensitivity function* capturing dependence on probabilities) and parametrizing risk aversion in the one-stage utility functions using a continuous real variable. By contrast, allowing the size of risk aversion shifts to depend on the magnitude of outcomes would be a more substantial change. Finding the history assignment involves a fixed point problem which would then become quite difficult to solve. The testable implications of such a model depend on whether it is possible to identify the extent to which a realization is disappointing or elating, as that designation depends on the extent to which other outcomes are considered elating or disappointing.

Finally, this paper considers a finite-horizon model of decision making. In an infinite-horizon setting, our methods extend to prove necessity of the reinforcement and primacy effects. However, our methods do not immediately extend to ensure the existence of an infinite-horizon internally consistent history assignment. One possible way to embed the finite-horizon HDRA preferences into an infinite-horizon economy is through the use of an overlapping generations model.

¹⁵An appendix regarding this extension will be provided upon request.

Appendix

Proof of Theorem 1. We first prove a sequence of four lemmas. The first two lemmas relate to necessity of the reinforcement and primacy effects. The last two lemmas relate to sufficiency.

Lemma 1. *Suppose an internally consistent history assignment exists. Then, for any h and t , and any h' with length t , we have $V_{hd^t} >_{RA} V_{hh'} >_{RA} V_{he^t}$.*

Proof. For simplicity and without loss of generality, we assume $h = 0$ because one can append the lotteries constructed below to a beginning lottery where each stage consists of getting a continuation lottery or a prize $z \in \{b, w\}$. In an abuse of notation, if we write a lottery or prize as an outcome when there are more stages left than present in the outcome, we mean that outcome is received for sure after the appropriate number of riskless stages (e.g., x instead of δ_x^ℓ). We proceed by induction. For $t = 1$, this is the reinforcement effect. Suppose it is not the case that $V_d >_{RA} V_e$. Since \mathcal{V} is ranked, this means that $V_e >_{RA} V_d$, or $CE_e(p) < CE_d(p)$ for any nondegenerate p . Pick any nondegenerate p and take $x \in (CE_e(p), CE_d(p))$. Then $\langle \alpha, p; 1 - \alpha, x \rangle$ has no internally consistent assignment. Now assume the claim holds for all $s \leq t - 1$, and suppose by contradiction that V_{d^t} is not the most risk averse. Then there is h' of length t such that $V_{h'} >_{RA} V_{d^t}$. It must be that $h' = eh''$ where h'' has length $t - 1$, otherwise there is a contradiction to the inductive step using $h = d$. By the inductive step, $V_{eh''}$ is less risk averse than $V_{ed^{t-1}}$, so $V_{ed^{t-1}} >_{RA} V_{d^t}$. Thus for any nondegenerate p , $CE_{ed^{t-1}}(p) < CE_{d^t}(p)$. Iteratively define the lottery p^{t-1} by $p^2 = \langle \alpha, p; 1 - \alpha, b \rangle$, and for each $3 \leq s \leq t - 1$, $p^s = \langle \alpha, p^{s-1}; 1 - \alpha, b \rangle$. Finally, let $p^t = \langle \beta, p^{t-1}; 1 - \beta, x \rangle$, where $x \in (CE_{ed^{t-1}}(p), CE_{d^t}(p))$. Note that the assignment of p must be d^{t-1} within p^{t-1} and that for α close to 1, the value of p^{t-1} is either close to $CE_{ed^{t-1}}(p)$ (if p^{t-1} is an elation) or close to $CE_{d^t}(p)$ (if p^{t-1} is a disappointment). But then for α close enough to 1, there is no consistent decomposition given the choice of x . Hence V_{d^t} is most risk averse. Analogously, to show that V_{e^t} is least risk averse, assuming it is not true implies $CE_{dh''}(p) > CE_{e^t}(p)$, and a similar construction with w instead of b in p^t , $x \in (CE_{e^t}(p), CE_{dh''}(p))$ and α close to 1, yields a contradiction. ■

Lemma 2. *Suppose an internally consistent history assignment exists. Then, for any h and t , we have $V_{hde^t} >_{RA} V_{hed^t}$.*

Proof. We use the same simplifications as the previous lemma (without loss of generality), and proceed by induction. For $t = 1$, suppose by contradiction that $V_{ed} >_{RA} V_{de}$. Consider a nondegenerate lottery p with $w \notin \text{supp } p$ (where supp denotes the support). For any $\beta \in (0, 1)$, let

$p^3 = \langle \beta, p^2, 1 - \beta, x \rangle$, where $p^2 = \langle \varepsilon(1 - \alpha), w; \varepsilon\alpha, p; 1 - 2\varepsilon, q \rangle$ and $\varepsilon, \alpha, q, x$ are chosen as follows. We want q to necessarily be elating (disappointing) if p^2 is disappointing (elating). Consider the conditions

$$\begin{aligned} CE_{dd}(q) &> \tau_d(\langle \alpha, CE_{de}(p); 1 - \alpha, w \rangle), \\ CE_{ee}(q) &< \min\{\tau_e(\langle \alpha, CE_{ee}(p); 1 - \alpha, w \rangle), CE_{ed}(p)\}. \end{aligned}$$

By Lemma 1, $\tau_d(\langle \alpha, CE_{de}(p); 1 - \alpha, w \rangle) < \tau_e(\langle \alpha, CE_{ee}(p); 1 - \alpha, w \rangle)$ for each choice of α, β . This is because the lottery on the RHS first-order stochastically dominates that on the LHS, and moreover is evaluated using a (weakly) less risk-averse threshold. By monotonicity of τ_h in α , choose α such that $\tau_e(\langle \alpha, CE_{ee}(p); 1 - \alpha, w \rangle) < CE_{ed}(p)$. Choose any nondegenerate q such that

$$\text{supp } q \subseteq \tau_d(\langle \alpha, CE_{de}(p); 1 - \alpha, \tau_e(\langle \alpha, CE_{ee}(p); 1 - \alpha, w \rangle) \rangle).$$

Using betweenness, this condition on q implies that it must be elating (disappointing) when p^2 is disappointing (elating). For ε sufficiently small, the value of p^2 is either close to $CE_{ed}(q)$ (when p^2 is elating) or close to $CE_{de}(q)$ (when p^2 is disappointing). Pick $x \in (CE_{ed}(q), CE_{de}(q))$ and notice that p^3 has no internally consistent history assignment.

Assume the lemma is true for $s \leq t - 1$. We prove it for $s = t$ by first proving $V_{de^t} >_{RA} V_{ed^{t-1}e}$. Suppose $V_{ed^{t-1}e} >_{RA} V_{de^t}$ by contradiction. Define, for any $p_1, \dots, p_{t-1} \in \mathcal{L}^1$, and $s \in \{2, \dots, t\}$,

$$\begin{aligned} a_s(\alpha) &:= \tau_{de^{t-s}}(\langle \alpha, CE_{de^{t-s}e}(p_{s-1}); 1 - \alpha, w \rangle), \\ b_s(\alpha) &:= \tau_{ed^{t-s}}(\langle \alpha, CE_{ed^{t-s}d}(p_{s-1}); 1 - \alpha, w \rangle). \end{aligned}$$

Pick \underline{y}, \bar{y} such that $w < \underline{y} < \bar{y} < b$ and ensure α is sufficiently small that $\tau_{ed^{t-1}}(\langle \alpha, \bar{y}; 1 - \alpha, w \rangle) < \underline{y}$. Take a nondegenerate $p_1 \in \mathcal{L}^1$ where $\{\bar{y}, \underline{y}\} \subseteq \text{supp } p_1 \subseteq [\underline{y}, \bar{y}]$. Let $a(\alpha) := a_1(\alpha)$. Now construct a sequence p_2, \dots, p_{t-1} where $\text{supp } p_s = \text{supp } p_1$ and $a_s(\alpha) = a$ for each s , as follows. To construct p_2 , compare $\tau_{de^{t-3}}(\langle \alpha, CE_{de^{t-3}d}(p_1); 1 - \alpha, w \rangle)$ with $\tau_{de^{t-2}}(\langle \alpha, CE_{de^{t-2}d}(p_1); 1 - \alpha, w \rangle)$. If the latter (resp., former) is the smaller of the two, construct p_2 by a first-order reduction (reps., improvement) in p_1 by mixing with $\delta_{\underline{y}}$ (reps., $\delta_{\bar{y}}$) using the appropriate weight, which exists because τ_h satisfies betweenness. Similarly construct the rest of the sequence. By the inductive step, $a_s(\alpha) < b_s(\alpha)$ for each $s \in \{2, \dots, t\}$ and the sequence of p_s above, $a_s(\alpha) = a(\alpha)$. Therefore,

$$\bigcap_{s=2}^t (a_s(\alpha), b_s(\alpha)) = (a(\alpha), \min_{s \in \{2, \dots, t\}} b_s(\alpha)) \neq \emptyset.$$

Let q be a nondegenerate lottery with $\text{supp } q \subseteq (a(\alpha), \min_{s \in \{2, \dots, t\}} b_s(\alpha))$. Construct the lottery $q^2 = \langle 1 - \varepsilon, q; \varepsilon, w \rangle$ and for each $s \in \{3, \dots, t+1\}$, define

$$q^s = \langle \varepsilon(1 - \alpha), w; \varepsilon\alpha, p_{s-2}; 1 - \varepsilon, q^{s-1} \rangle.$$

Finally, define $q^{t+2} = \langle \gamma, q^{t+1}; 1 - \gamma, x \rangle$ where $x \in (CE_{ed^{t-1}e}(q), CE_{de^t}(q))$. Using Lemma 1 and the choice of q 's support in the interval above, each sublottery q^s (for $2 \leq s \leq t$) is disappointing (elating) if q^{t+1} is elating (disappointing). For ε sufficiently small, the value of q^{t+1} is either very close to $CE_{ed^{t-1}e}(q)$ when it is elating or $CE_{de^t}(q)$ when it is disappointing. But by the choice of x , there is no internally consistent history assignment.

To complete the proof, assume by contradiction that $V_{ed^t} >_{RA} V_{de^t}$. Recall a_s, b_s and the sequence p_1, \dots, p_{t-1} , constructed so $a_s(\alpha) = a(\alpha)$ for every $s \in \{2, \dots, t\}$. Define, for any p_0 ,

$$\begin{aligned} a_1(\alpha) &:= \tau_{de^{t-1}}(\langle \alpha, CE_{de^{t-1}e}(p_0); 1 - \alpha, w \rangle), \\ b_1(\alpha) &:= \tau_{ed^{t-1}}(\langle \alpha, CE_{ed^{t-1}e}(p_0); 1 - \alpha, w \rangle). \end{aligned}$$

Notice that $a_1 < b_1$ by the claim we just proved and the inductive hypothesis applied to τ_h (that is, $\tau_{de^{t-1}}$ is weakly more risk averse than $\tau_{ed^{t-1}}$). Construct p_0 with the same support as p_1 such that $a_1(\alpha) = a$. By the choice of α , notice that $b_1 < CE_{ed^{t-1}d}(p_0)$. Let $\tilde{q} \in \mathcal{L}^1$ be nondegenerate, with $\text{supp } \tilde{q} \subseteq (a(\alpha), \min_{s \in \{1, \dots, t\}} b_s(\alpha))$. For each $s \in \{2, \dots, t+1\}$, define

$$\tilde{q}^s = \langle \varepsilon(1 - \alpha), w; \varepsilon\alpha, p_{s-2}; 1 - \varepsilon, \tilde{q}^{s-1} \rangle.$$

Finally, define $\tilde{q}^{t+2} = \langle \gamma, \tilde{q}^{t+1}; 1 - \gamma, x \rangle$ where $x \in (CE_{ed^{t-1}e}(\tilde{q}), CE_{de^t}(\tilde{q}))$. For ε sufficiently small, the certainty equivalent of \tilde{q}^{t+1} is either very close to $CE_{ed^{t-1}e}(\tilde{q})$ when it is elating or $CE_{de^t}(\tilde{q})$ when it is disappointing. But then \tilde{q}^{t+1} has no internally consistent assignment. ■

Lemma 3. For any $V : \mathcal{L}^1 \rightarrow \mathbb{R}$, define, for any $p \in \mathcal{L}^1$, $e(p) := \{x \in \text{supp } p \mid V(\delta_x) > V(p)\}$ and $d(p) := \{x \in \text{supp } p \mid V(\delta_x) < V(p)\}$. Take lotteries $p = \langle \alpha_1, x_1; \alpha_2, x_2; \dots; \alpha_m, x_m \rangle$ and $p' = \langle \alpha_1, x'_1; \alpha_2, x_2; \dots; \alpha_m, x_m \rangle$ with the same $(\alpha_i)_{i=1}^m$ and $(x_i)_{i=2}^m$, but $x_1 \neq x'_1$. For V satisfying betweenness: (1) if $x_1 \notin d(p)$ and $x'_1 > x_1$ then $x'_1 \in e(p')$; and (2) if $x_1 \notin e(p)$ and $x'_1 < x_1$ then $x'_1 \in d(p')$.

Proof. We prove statement (1), since the proof of (2) is analogous. If $x_1 \notin d(p)$ then $\delta_{x_1} \succeq p$, where \succeq represents V . Note that p can be written as the convex combination of the lotteries δ_{x_1} (with weight α_1) and $p_{-1} = \langle \frac{\alpha_2}{1-\alpha_1}, x_2; \dots; \frac{\alpha_m}{1-\alpha_1}, x_m \rangle$ (with weight $1 - \alpha_1$). By betweenness, $\delta_{x_1} \succeq p_{-1}$. Since $x'_1 > x_1$, monotonicity implies $\delta_{x'_1} \succ \delta_{x_1} \succeq p_{-1}$, and thus that $\delta_{x'_1} \succ p_{-1}$. But then again by

betweenness and the fact that $\alpha_1 \in (0, 1)$, x'_1 is strictly preferred to the convex combination of $\delta_{x'_1}$ (with weight α_1) and p_{-1} (with weight $1 - \alpha_1$), which is simply p' . Hence $x'_1 \in e(p')$. ■

Lemma 4. *Consider $T = 2$, any (endogenous or exogenous) threshold rule, and suppose that $V_d >_{RA} V_e$. Then any nondegenerate $p^2 \in \mathcal{L}^2$ has an internally consistent history assignment (using only strict elation and disappointment for nondegenerate lotteries in its support).*

Proof. Consider $p^2 = \langle \alpha_1, p_1; \dots; \alpha_m, p_m \rangle$. Suppose for simplicity that all p_i are nondegenerate (if $p_i = \delta_x$ is degenerate, then $CE_e(p_i) = CE_d(p_i)$, so the algorithm can be run on the nondegenerate sublotteries, with the degenerate ones labeled ex-post according to internal consistency). Without loss of generality, suppose that the indexing in p^2 is such that $p_1 \in \arg \max_{i=1, \dots, m} CE_e(p_i)$, $p_m \in \arg \min_{i=2, \dots, m} CE_d(p_i)$, and $CE_e(p_2) \geq CE_e(p_3) \geq \dots \geq CE_e(p_{m-1})$. Initialize the algorithm by $a^1(p_1|p^2) = e$ and $a^1(p_j|p^2) = d$ for all $i > 1$. Let τ^1 be the threshold value of p^2 when it is folded back under a^1 ; if the assignment of p^2 under a^1 is consistent given τ^1 , the algorithm and proof are complete. If not, consider $i = 2$. If $CE_d(p_2) \geq \tau^1$, then set $a^2(p_2|p^2) = e$ and $a^2(p_i|p^2) = a^1(p_i|p^2)$ for all $i \neq 2$ (if $CE_d(p_2) < \tau^1$, let $a^2(p_i|p^2) = a^1(p_i|p^2)$ for all i). Let τ^2 be the resulting threshold value of p^2 when it is folded back under a^2 . If the assignment of p^2 under a^2 is consistent given τ^2 , the algorithm and proof are complete. If not, move to $i = 3$, and so on, so long as $i \leq m - 1$. Since the threshold rule is from the betweenness class, Lemma 3 implies that if $CE_d(p_i) \geq \tau^{i-1}$, then $CE_e(p_i) > \tau^i$. Moreover, notice that if $CE_e(p_i) > \tau^i$, then for any $j < i$, $CE_e(p_j) \geq CE_e(p_i) > \tau^i$, so previously switched assignments remain strict elations; also, because $\tau^i \geq \tau^{i-1}$ for all i , previous disappointments remain disappointments. If the final step of the algorithm reaches $i = m - 1$, notice that $CE_d(p_m)$ is the smallest disappointment certainty equivalent, hence the smallest in $\{CE_{a^{m-1}(p_j)}(p_j)\}_{j=1, \dots, m}$. The final assignment of p^2 under a^{m-1} is thus consistent given τ^{m-1} . ■

We now complete the proof of Theorem 1. The reinforcement and primacy effects are necessary by Lemmas 1 and 2. By Lemma 4 and the reinforcement effect, an internally consistent (strict) history assignment exists for any nondegenerate $p \in \mathcal{L}^2$, using any initial V_h . By induction, suppose that for any $(t - 1)$ -stage lottery an internally consistent history assignment exists, using any initial V_h . Now consider a sublottery $p^t = \langle \alpha_1, p_1^{t-1}; \dots; \alpha_m, p_m^{t-1} \rangle$. The algorithm in Lemma 4 for \mathcal{L}^2 only requires $CE_e(p) > CE_d(p)$ for any nondegenerate $p \in \mathcal{L}^1$. The same algorithm constructs an internally consistent assignment within p^t if for any $p^{t-1} \in \mathcal{L}^{t-1}$ such that $p^{t-1} \neq \delta_x^{t-1}$ for some x , the recursive certainty equivalent of p^{t-1} is strictly higher under elation than disappointment. While there may be multiple consistent assignments within p^{t-1} , the primacy and reinforcement effects ensure this strict comparison given any history assignment. Starting with V_{he} , the lottery is folded back using higher certainty equivalents per sublottery, and using a (weakly) less risk averse

threshold, as compared to starting with the more risk averse V_{hd} . In the case $p^{t-1} = \delta_x^{t-1}$ for some x , the assignment can be made ex-post according to what is consistent. ■

Proof of Proposition 1. Rearrange Equation (6) to

$$\frac{u_h\left(u_{he}^{-1}\left((1-\alpha)u_{he}(H)+\alpha u_{he}(L)\right)\right)-u_h(L)}{1-\alpha} > \frac{u_h\left(u_{he}^{-1}\left(\alpha u_{he}(H)+(1-\alpha)u_{he}(L)\right)\right)-u_h(L)}{\alpha}.$$

We may write $u_h = f \circ u_{he}$ for some f ; hence $u_h^{-1} = u_{he}^{-1} \circ f$. Dividing by $u_{he}(H) - u_{he}(L) > 0$,

$$\frac{f\left(u_{he}(L)+(1-\alpha)(u_{he}(H)-u_{he}(L))\right)-f(u_{he}(L))}{(1-\alpha)(u_{he}(H)-u_{he}(L))} > \frac{f\left(u_{he}(L)+\alpha(u_{he}(H)-u_{he}(L))\right)-f(u_{he}(L))}{\alpha(u_{he}(H)-u_{he}(L))}.$$

These are slopes of segments joining f 's graph. As $\alpha > .5 \Leftrightarrow \alpha(u_{he}(H)-u_{he}(L)) > (1-\alpha)(u_{he}(H)-u_{he}(L))$, the above holds for all $\alpha \in (.5, 1)$ and $H > L$ if and only if f is concave, or $V_{he} <_{RA} V_h$. ■

Threshold elicitation. We define \succeq_{τ_h} based on the DM's preference \succeq over \mathcal{L}^T as follows. For $x \in X$ and $q \in \mathcal{L}^1$, we say $\delta_x \succ_{\tau} q$ (resp., $q \succ_{\tau} \delta_x$) if there is $x' < x$ (resp., $x' > x$) and a nondegenerate $p \in \mathcal{L}^1$ such that $\ell_{he}(\delta_{x'}) \sim \ell_{he}(p)$ (resp., $\ell_{hd}(\delta_{x'}) \sim \ell_{hd}(p)$) and $m_h(p, q) \sim m_h(\delta_{x'}, q)$. For any $p, q \in \mathcal{L}^1$, we say $p \succ_{\tau_h} q$ if there is $x \in X$ such that $p \succ_{\tau_h} \delta_x \succ_{\tau_h} q$. We say $p \sim_{\tau_h} q$ if $p \not\succeq_{\tau_h} q$ and $q \not\succeq_{\tau_h} p$.

Proposition 2. $\tau_h(p) > \tau_h(q)$ if and only if $p \succ_{\tau_h} q$.

Proof. If $\tau_h(p) > \tau_h(q)$, then take $x' > x > x''$ such that $\tau_h(p) > \tau_h(\delta_{x'}) = x' > \tau_h(\delta_x) = x > \tau_h(\delta_{x''}) = x'' > \tau_h(q)$. By the assumptions on V_h , we can pick $\alpha', \alpha'' \in (0, 1)$ such that $r' = \langle \alpha', x'; 1 - \alpha', x'' \rangle$ and $r'' = \langle \alpha'', x'; 1 - \alpha'', x'' \rangle$ satisfy $CE_{hd}(r') = CE_{he}(r'') = x$. Thus, $\ell_{hd}(r') \sim \ell_{hd}(\delta_x)$ and $\ell_{he}(r'') \sim \ell_{he}(\delta_x)$. Either $m_h(r', p) \sim m_h(\delta_x, p)$ or $m_h(r', p) \sim m_h(\delta_{CE_{he}(r')}, p)$ must hold, and the latter is impossible because $x < x' < \tau_h(p)$ implies $CE_{he}(r') < x' < \tau_h(p)$. Thus, by definition, $p \succ_{\tau_h} \delta_x$. An analogous argument proves $\delta_x \succ_{\tau_h} \delta_q$, which means that $p \succ_{\tau_h} q$.

Now, suppose $p \succ_{\tau_h} q$. Hence there is x such that $p \succ_{\tau_h} \delta_x \succ_{\tau_h} q$, and so there exists $x' > x$ such that $p \succ_{\tau_h} \delta_{x'}$ and $x'' < x$ such that $\delta_{x''} \succ_{\tau_h} q$. Since $m_h(r, p) \sim m_h(\delta_{x'}, p)$ for some nondegenerate r such that $CE_{hd}(r) = x'$, internal consistency implies $\tau_h(\delta_{x'}) < \tau_h(\langle .5, \delta_{x'}; .5p(x_1), \delta_{x_1}; \dots; .5p(x_n), \delta_{x_n} \rangle)$, which by betweenness means that $\tau_h(\delta_{x'}) < \tau_h(p)$. Similarly, from $\delta_{x''} \succ_{\tau_h} q$, we have $\tau_h(\delta_{x''}) > \tau_h(q)$. Since $x'' < x'$, this concludes the proof. ■

Proof of Theorem 2. The proof of necessity is analogous to that of Theorem 1. The proof of sufficiency is analogous as well, with two additions of note. First, since the reinforcement and

strong primacy effects imply the certainty equivalent of each decision problem in a choice set increases when evaluated as an elation, the certainty equivalent of the choice set (the maximum of those values) also increases when viewed as an elation (relative to being viewed as a disappointment). Second, if the certainty equivalent of a choice set is the same when viewed as an elation and as a disappointment, the best option in both choice sets must be degenerate. Then its history assignment may be made ex-post according to internal consistency. ■

Proof of Theorem 3. The proof proceeds in steps.

Step 1: The agent's optimization problem given a history assignment

We solve the agent's problem by backward induction, denoting by $h(y_1, \dots, y_{t-1})$ the agent's history assignment after the sequence of realizations (y_1, \dots, y_{t-1}) and later solving for the right assignment. After the realizations (y_1, y_2) but before learning y_3 , the agent solves the problem:

$$\max_{\alpha(y_1, y_2)} \mathbb{E}_{\tilde{y}_3} \left[u_{h(y_1, y_2)} \left(c(y_1) + c(y_2|y_1) + \alpha(y_1, y_2)\tilde{y}_3 + (1 - \alpha(y_1, y_2))P(y_1, y_2)R \right) \right],$$

where $c(y_1) = \alpha y_1 + (1 - \alpha)PR$ denotes the realized consumption in period 1 and $c(y_2|y_1) = \alpha(y_1)y_2 + (1 - \alpha(y_1))P(y_1)R$ denotes the realized consumption in period 2. Using the CARA form, the first-order condition given y_1 and y_2 simplifies to

$$P(y_1, y_2)R = \frac{1}{1 + \exp(\lambda_{h(y_1, y_2)} \alpha(y_1, y_2)(H - L))} H + \frac{\exp(\lambda_{h(y_1, y_2)} \alpha(y_1, y_2)(H - L))}{1 + \exp(\lambda_{h(y_1, y_2)} \alpha(y_1, y_2)(H - L))} L. \quad (9)$$

Let $\alpha^*(y_1, y_2)$ be the optimal choice and let $c^*(\tilde{y}_3|y_1, y_2) = \alpha^*(y_1, y_2)\tilde{y}_3 + (1 - \alpha^*(y_1, y_2))P(y_1, y_2)R$ denote the optimal consumption plan for $t = 3$ given y_1 and y_2 . Using the notation $\Gamma_{\tilde{x}}(\lambda, \phi(\tilde{x}))$ from the text, and the CARA functional form, the recursive certainty equivalent $RCE(y_1, y_2)$ of the choice problem in $t = 3$ after the realizations (y_1, y_2) is then:

$$RCE(y_1, y_2) = c(y_1) + c(y_2|y_1) + \Gamma_{\tilde{y}_3} \left(\lambda_{h(y_1, y_2)}, c^*(\tilde{y}_3|y_1, y_2) \right),$$

which is the sum of period-1 consumption, period-2 consumption, and the certainty equivalent of period-3 optimal consumption given risk aversion $\lambda_{h(y_1, y_2)}$. Proceeding backwards, after observing

y_1 but before learning y_2 , the agent solves the problem¹⁶

$$\max_{\alpha(y_1)} \mathbb{E}_{\tilde{y}_2} \left[u_{h(y_1)} \left(RCE(y_1, \tilde{y}_2) \right) \right],$$

where $RCE(y_1, y_2)$, defined above, is a function of $\alpha(y_1)$ through $c(y_2|y_1)$. Using the CARA form, the first-order condition at $t = 2$ simplifies to

$$P(y_1)R = \frac{1}{1 + \exp(\lambda_{h(y_1)}(RCE(y_1, H) - RCE(y_1, L)))} H + \frac{\exp(\lambda_{h(y_1)}(RCE(y_1, H) - RCE(y_1, L)))}{1 + \exp(\lambda_{h(y_1)}(RCE(y_1, H) - RCE(y_1, L)))} L. \quad (10)$$

Let $\alpha^*(y_1)$ be the optimal choice and denote by $c^*(\tilde{y}_2|y_1) = \alpha^*(y_1)\tilde{y}_2 + (1 - \alpha^*(y_1))P(y_1)R$ the agent's optimal consumption plan for $t = 2$ given y_1 . The recursive certainty equivalent $RCE(y_1)$ of the choice problem in $t = 2$ after the realization y_1 is then:

$$RCE(y_1) = c(y_1) + \Gamma_{\tilde{y}_2} \left(\lambda_{h(y_1)}, c^*(\tilde{y}_2|y_1) + \Gamma_{\tilde{y}_3} \left(\lambda_{h(y_1, \tilde{y}_2)}, c^*(\tilde{y}_3|y_1, \tilde{y}_2) \right) \right).$$

The random variable $c^*(\tilde{y}_2|y_1) + \Gamma_{\tilde{y}_3} \left(\lambda_{h(y_1, \tilde{y}_2)}, c^*(\tilde{y}_3|y_1, \tilde{y}_2) \right)$, which is a function of \tilde{y}_2 , is the sum of period-2 optimal consumption and the certainty equivalent of period-3 optimal consumption given risk aversion $\lambda_{h(y_1, \tilde{y}_2)}$. The term $RCE(y_1)$ is simply the certainty equivalent of this random variable given risk aversion $\lambda_{h(y_1)}$.

Should the agent choose to hold an initial share α of the risky asset, the value of the decision problem the agent faces at $t = 1$ is given by $\mathbb{E}_{\tilde{y}_1} [u_0(RCE(\tilde{y}_1))]$, where $RCE(y_1)$ depends on α through $c(y_1)$. The agent then chooses the optimal initial share α^* :

$$\max_{\alpha} \mathbb{E}_{\tilde{y}_1} \left[u_0 \left(RCE(\tilde{y}_1) \right) \right].$$

Using the CARA form, and letting P be the initial price of the risky asset in period $t = 1$ prior to any realizations, the first-order condition at $t = 1$ simplifies to

$$PR = \frac{1}{1 + \exp(\lambda_0(RCE(H) - RCE(L)))} H + \frac{\exp(\lambda_0(RCE(H) - RCE(L)))}{1 + \exp(\lambda_0(RCE(H) - RCE(L)))} L. \quad (11)$$

¹⁶Note that the $t = 1, 2$ problems fix future histories regardless of the choice of α and $\alpha(y_1)$. Because period-3 prices will be such that the agent holds the asset at *some* positive coefficient of risk aversion given (y_1, y_2) , we know that $P(y_1, y_2) < \frac{1}{R} \frac{H+L}{2}$. In period 3, the agent is thus willing to hold some amount of the asset at *any* level of risk aversion. This means that in periods $t = 1, 2$, even if the agent were hypothetically not to hold the asset, he would still be exposed to risk at $t = 1, 2$ due to the influence on period-3 prices; hence the agent's history is as specified.

In view of (9), (10), and (11) and the fact that $\exp(\cdot) \geq 0$, we conclude that all the $P(y_1, \dots, y_{t-1})R$'s are a convex combination of H and L , where the weights depend on the agent's risk aversion after the realizations y_1, \dots, y_{t-1} . To determine those levels of risk aversion, we use the equilibrium condition that the representative agent must optimally hold his per-period endowment of one share of risky asset after any realization and given his history assignment. That is, we plug $\alpha^*(y_1, y_2) = \alpha^*(y_1) = \alpha^* = 1$ into the formulas for $RCE(y_1, y_2)$ and $RCE(y_1)$ and look for an internally consistent history assignment in order to deduce the equilibrium prices.

Step 2: Verifying internal consistency of the history assignment

We now check that it is internally consistent for the agent to consider each *High* realization elating and each *Low* realization disappointing (Lemma 5 below shows that this is also the unique internally consistent history assignment). We proceed recursively, fixing a realization $y_1 \in \{L, H\}$. Being elated by $y_2 = H$ and disappointed by $y_2 = L$ is internally consistent if the resulting certainty equivalents for the $t = 3$ choice problems satisfy

$$RCE(y_1, H) \geq \tau_{h(y_1)} \left(\left\langle \frac{1}{2}, RCE(y_1, H); \frac{1}{2}, RCE(y_1, L) \right\rangle \right) > RCE(y_1, L).$$

But this holds if and only if $RCE(y_1, H) > RCE(y_1, L)$, which in turns holds if and only if

$$H + \Gamma_{\tilde{y}_3}(\lambda_{h(y_1)e}, \tilde{y}_3) > L + \Gamma_{\tilde{y}_3}(\lambda_{h(y_1)d}, \tilde{y}_3). \quad (12)$$

As λ increases, risk aversion increases and the certainty equivalent $\mu_{\tilde{y}_3}(\lambda, \tilde{y}_3)$ decreases. Hence the reinforcement effect, or $\lambda_{h(y_1)e} < \lambda_{h(y_1)d}$, implies that Equation (12) must hold.

Proceeding backwards, observe that being elated by $y_1 = H$ and disappointed by $y_2 = L$ is internally consistent if, similarly to our previous calculation, the resulting recursive certainty equivalents for the $t = 2$ choice problems satisfy $RCE(H) > RCE(L)$. In turn, this holds if and only if

$$H + \Gamma_{\tilde{y}_2} \left(\lambda_e, \tilde{y}_2 + \Gamma_{\tilde{y}_3}(\lambda_{h(H, \tilde{y}_2)}, \tilde{y}_3) \right) > L + \Gamma_{\tilde{y}_2} \left(\lambda_d, \tilde{y}_2 + \Gamma_{\tilde{y}_3}(\lambda_{h(L, \tilde{y}_2)}, \tilde{y}_3) \right), \quad (13)$$

where the history assignments above are $h(H, H) = ee$, $h(H, L) = ed$, $h(L, H) = de$, and $h(L, L) = dd$. Using this, it is easy to compare the certainty equivalent $\Gamma_{\tilde{y}_3}(\lambda_{h(y_1, \tilde{y}_2)}, \tilde{y}_3)$ on each side of Equation (13). By the primacy and reinforcement effects, $\lambda_{ee} < \lambda_{ed} < \lambda_{de} < \lambda_{dd}$. Thus, given any realization of \tilde{y}_2 , the random variable $\tilde{y}_2 + \Gamma_{\tilde{y}_3}(\lambda_{h(y_1, \tilde{y}_2)}, \tilde{y}_3)$ takes a larger value when $y_1 = H$ than when $y_1 = L$. Moreover, due to the reinforcement effect, the random variable is also evaluated using a less risk-averse coefficient when $y_1 = H$ than when $y_1 = L$. Hence Equation (13) must also

hold, and our proposed history assignment is internally consistent. The next lemma shows that it is the unique internally consistent assignment.

Lemma 5. *The history assignment is unique.*

Proof. If it were internally consistent after y_1 for H to be disappointing and L to be elating, that would mean $L + \Gamma_{\tilde{y}_2}(\lambda_{h(y_1)e}, \tilde{y}_2) > H + \Gamma_{\tilde{y}_2}(\lambda_{h(y_1)d}, \tilde{y}_2)$, or $H - L < \Gamma_{\tilde{y}_2}(\lambda_{h(y_1)e}, \tilde{y}_2) - \Gamma_{\tilde{y}_2}(\lambda_{h(y_1)d}, \tilde{y}_2)$. Notice that $\Gamma_{\tilde{y}_2}(\lambda_{h(y_1)e}, \tilde{y}_2) < \frac{1}{2}H + \frac{1}{2}L$ (the boundary case $\lambda_{h(y_1)e} = 0$) and $\Gamma_{\tilde{y}_2}(\lambda_{h(y_1)d}, \tilde{y}_2) > L$ (the case $\lambda_{h(y_1)d} = \infty$). Thus, $H - L < \Gamma_{\tilde{y}_2}(\lambda_{h(y_1)e}, \tilde{y}_2) - \Gamma_{\tilde{y}_2}(\lambda_{h(y_1)d}, \tilde{y}_2) < \frac{1}{2}H - \frac{1}{2}L$, a contradiction.

Suppose by contradiction that in the first period, H is disappointing and L is elating. Internal consistency requires $L + \Gamma_{\tilde{y}_2}(\lambda_e, \tilde{y}_2 + \Gamma_{\tilde{y}_3}(\lambda_{h(L, \tilde{y}_2)}, \tilde{y}_3)) > H + \Gamma_{\tilde{y}_2}(\lambda_d, \tilde{y}_2 + \Gamma_{\tilde{y}_3}(\lambda_{h(H, \tilde{y}_2)}, \tilde{y}_3))$. But since this is also the certainty equivalent of a (more complex) random variable, we have the bound $\Gamma_{\tilde{y}_2}(\lambda_d, \tilde{y}_2 + \Gamma_{\tilde{y}_3}(\lambda_{h(H, \tilde{y}_2)}, \tilde{y}_3)) > L + \Gamma_{\tilde{y}_3}(\lambda_{dd}, \tilde{y}_3) > 2L$. Similarly, we also have the bound $\Gamma_{\tilde{y}_2}(\lambda_e, \tilde{y}_2 + \Gamma_{\tilde{y}_3}(\lambda_{h(L, \tilde{y}_2)}, \tilde{y}_3)) < \frac{1}{2}(H + L + \Gamma_{\tilde{y}_3}(\lambda_{ee}, \tilde{y}_3) + \Gamma_{\tilde{y}_3}(\lambda_{ed}, \tilde{y}_3)) < H + L$. But then we have a contradiction $H - L < \Gamma_{\tilde{y}_2}(\lambda_e, \tilde{y}_2 + \Gamma_{\tilde{y}_3}(\lambda_{h(L, \tilde{y}_2)}, \tilde{y}_3)) - \Gamma_{\tilde{y}_2}(\lambda_d, \tilde{y}_2 + \Gamma_{\tilde{y}_3}(\lambda_{h(H, \tilde{y}_2)}, \tilde{y}_3)) < H - L$. ■

Step 3: Implications for prices

Given Steps 1 and 2, the asset price in period t depends on the sequence of realizations y_1, \dots, y_{t-1} and is given by Equation (8) and Table 1. Given the discussion in the text of how the price comparisons follow, it remains to prove two lemmas.

Lemma 6. $P(H) > P(L)$.

Proof. Using (10), we know $P(H) > P(L)$ if and only if the following function is negative:

$$G(\lambda_0, H, L) := \lambda_0(H - L)(a^2 - b^2) + \lambda_0 \left(a^2 (\Gamma_{\tilde{y}_3}(a^3 \lambda_0, \tilde{y}_3) - \Gamma_{\tilde{y}_3}(a^2 b \lambda_0, \tilde{y}_3)) - b^2 (\Gamma_{\tilde{y}_3}(b^2 a \lambda_0, \tilde{y}_3) - \Gamma_{\tilde{y}_3}(b^3 \lambda_0, \tilde{y}_3)) \right),$$

where \tilde{y}_3 is H or L with probability $1/2$. By pulling out the term $\exp(-\lambda L)$, note that $\Gamma_{\tilde{y}_3}(\lambda, \tilde{y}_3) = L - \frac{1}{\lambda} \ln \left(\frac{1}{2} \exp(-\lambda(H - L)) + \frac{1}{2} \right)$. Hence

$$\begin{aligned} G(\lambda_0, H, L) &= \lambda_0(H - L)(a^2 - b^2) + \frac{1}{b} \ln \left(\frac{1}{2} \exp(-\lambda_0 a^2 b(H - L)) + \frac{1}{2} \right) \\ &\quad - \frac{1}{a} \ln \left(\frac{1}{2} \exp(-\lambda_0 a^3(H - L)) + \frac{1}{2} \right) - \frac{1}{b} \ln \left(\frac{1}{2} \exp(-\lambda_0 b^3(H - L)) + \frac{1}{2} \right) \\ &\quad + \frac{1}{a} \ln \left(\frac{1}{2} \exp(-\lambda_0 a b^2(H - L)) + \frac{1}{2} \right). \end{aligned}$$

We want to show that $\frac{\partial G}{\partial H}(\lambda_0, H, L) < 0$, implying $G(\lambda_0, H, L)$ would be maximized at $H = L$, where it has value zero. The derivative of G with respect to H is given by

$$\frac{\partial G}{\partial H}(\lambda_0, H, L) = \lambda_0 \left(a^2(1 + g(a^3) - g(a^2b)) - b^2(1 - g(b^3) + g(ab^2)) \right),$$

using the definition $g(x) = \frac{\exp(-\lambda_0 x(H-L))}{\exp(-\lambda_0 x(H-L)) + 1}$. Note that $g'(x) = \frac{-\lambda_0 x(H-L) \exp(\lambda_0 x(H-L))}{x(1 + \exp(\lambda_0 x(H-L)))^2}$. Moreover, $-\frac{1}{x} < g'(x) < 0$ for all $x \geq 0$. Negativity of $g'(x)$ is clear. To see the left bound, simply observe that $\lambda_0 x(H-L) \exp(\lambda_0 x(H-L)) \leq (\exp(\lambda_0 x(H-L)))^2$. The mean value theorem says that for some $c_1 \in (a^3, a^2b)$ we have $g(a^3) - g(a^2b) = a^2(a-b)g'(c_1)$. Similarly, for some $c_2 \in (ab^2, b^3)$ we have $g(ab^2) - g(b^3) = b^2(a-b)g'(c_2)$. Then,

$$\frac{\partial G}{\partial H}(\lambda_0, H, L) = \lambda_0 \left(a^2 - b^2 + a^4(a-b)g'(c_1) - b^4(a-b)g'(c_2) \right) < \lambda_0 \left(a^2 - b^2 + a^4(a-b) \left(-\frac{1}{ab^2} \right) \right),$$

which equals $\lambda_0(a-b)(a+b - \frac{a^3}{b^2})$. This is negative as desired, since $0 < a < 1 < b$. ■

Lemma 7. $P(H) > P$.

Proof. Since $a < 1$, it suffices to show that $a^2(\Gamma_{\tilde{y}_3}(a^3\lambda_0, \tilde{y}_3) - \Gamma_{\tilde{y}_3}(a^2b\lambda_0, \tilde{y}_3))$ is smaller than $\Gamma_{\tilde{y}_2}(a^2\lambda_0, \tilde{y}_2 + \Gamma_{\tilde{y}_3}(\lambda_{h(H, \tilde{y}_2)}, \tilde{y}_3)) - \Gamma_{\tilde{y}_2}(b^2\lambda_0, \tilde{y}_2 + \Gamma_{\tilde{y}_3}(\lambda_{h(L, \tilde{y}_2)}, \tilde{y}_3))$. To show this,¹⁷ define $\hat{\Gamma}(\lambda, m, n) \equiv \Gamma_{\tilde{x}}(\lambda, \tilde{x})$ for the random variable \tilde{x} which gives each of m and n with probability one-half. Also, we define the notation $\mu = \hat{\Gamma}(a^3\lambda_0, H, L)$ and $\nu = \hat{\Gamma}(a^2b\lambda_0, H, L)$. Because $\hat{\Gamma}(\lambda, m, n)$ is a certainty equivalent, it is increasing in both m, n and decreasing in λ . Because $b^2\lambda_0 > ab$, $b^2a > a^3$, and $b^3 > a^2b$, it suffices to show $\hat{\Gamma}(a^2\lambda_0, H + \mu, L + \nu) - \hat{\Gamma}(ab\lambda_0, H + \mu, L + \nu) > a^2(\mu - \nu)$, which is stronger. Letting $\gamma(\lambda, m, n) = -\lambda^2 \frac{\partial \hat{\Gamma}}{\partial \lambda}(\lambda, m, n)$, we note γ satisfies three properties: (i) $\frac{\partial \gamma}{\partial \lambda} > 0$; (ii) if $m > n$, $\frac{\partial \gamma}{\partial m} > 0$; and (iii) $\gamma(\lambda, m + c, n + c) = \gamma(\lambda, m, n)$ for all $c > 0$. To see this, observe that

$$\gamma(\lambda, m, n) = \frac{-\lambda m \exp(-\lambda m) - \lambda n \exp(-\lambda n)}{\exp(-\lambda m) + \exp(-\lambda n)} - \ln\left(\frac{\exp(-\lambda m) + \exp(-\lambda n)}{2}\right).$$

Property (iii) then follows from simple algebra. Property (i) follows from the fact that

$$\frac{\partial \gamma}{\partial \lambda}(\lambda, m, n) = \frac{\lambda(m-n)^2 \exp(-\lambda(m+n))}{(\exp(-\lambda m) + \exp(-\lambda n))^2} > 0.$$

Using $p = \exp(-\lambda m)$, $q = \exp(-\lambda n)$ in γ , observe that the derivative of $\frac{p \ln p + q \ln q}{p+q} - \ln\left(\frac{p+q}{2}\right)$ with

¹⁷We thank Xiaosheng Mu for providing the argument showing this inequality.

respect to p is $\frac{q(\ln p - \ln q)}{(p+q)^2} < 0$, since $p < q$. Property (ii) follows by $\frac{dp}{dm} < 0$. To complete the proof,

$$\begin{aligned}
\hat{\Gamma}(a^2\lambda_0, H + \mu, L + \nu) - \Gamma(ab\lambda, H + \mu, L + \nu) &= \int_{a^2\lambda_0}^{ab\lambda_0} -\frac{\partial \hat{\Gamma}}{\partial \lambda}(x, H + \mu, L + \nu) dx \\
&\stackrel{\text{(by definition)}}{=} \int_{a^2\lambda_0}^{ab\lambda_0} \frac{1}{x^2} \gamma(x, H + \mu, L + \nu) dx \\
&\stackrel{\text{(by property (iii) of } \gamma)}{=} \int_{a^2\lambda_0}^{ab\lambda_0} \frac{1}{x^2} \gamma(x, H + \mu - \nu, L) dx \\
&\stackrel{\text{(by } \mu > \nu, a < 1, \text{ and properties (i)-(ii) of } \gamma)}{>} \int_{a^2\lambda_0}^{ab\lambda_0} \frac{1}{x^2} \gamma(ax, H, L) dx \\
&\stackrel{\text{(changing variables to } y = ax)}{=} \int_{a^3\lambda_0}^{a^2b\lambda_0} \frac{1}{(\frac{y}{a})^2} \gamma(y, H, L) \frac{1}{a} dy \\
&\stackrel{\text{(by definition)}}{=} a \int_{a^3\lambda_0}^{a^2b\lambda_0} -\frac{\partial \hat{\Gamma}}{\partial \lambda}(y, H, L) dy \\
&\stackrel{\text{(since } a < 1)}{>} a^2 (\hat{\Gamma}(a^3\lambda_0, H, L) - \hat{\Gamma}(a^2b\lambda_0, H, L)). \quad \blacksquare
\end{aligned}$$

This completes the proof of Theorem 3. ■

References

- Alesina, Alberto and Nicola Fuchs-Schündeln, “Good Bye Lenin (Or Not?): The Effect of Communism on People’s Preferences,” *American Economic Review*, 2007, 97, 1507–1528.
- Anderson, Norman, “Primacy effects in personality impression formation using a generalized order effect paradigm,” *Journal of Personality and Social Psychology*, 1965, 2, 1–9.
- Ang, Andrew, Geert Bekaert and Jun Liu, “Why Stocks May Disappoint,” *Journal of Financial Economics*, 2005, 76, 471–508.
- Baird, Jane and Robert Zelin II, “The Effects of Information Ordering on Investor Perceptions: An Experiment Utilizing Presidents Letters,” *Journal of Financial and Strategic Decisions*, 2000, 13(3), 71–80.
- Barberis, Nicholas, Ming Huang and Tano Santos, “Prospect Theory and Asset Prices,” *Quarterly Journal of Economics*, 2001, 116, 1–53.

- Cameron, Lisa and Manisha Shah, “Do Natural Disasters Shape Risk Attitudes?,” *mimeo*, 2010.
- Campbell, John and John Cochrane, “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 1999, CVII, 205–251.
- Caplin, Andrew and John Leahy, “Psychological Expected Utility Theory And Anticipatory Feelings,” *Quarterly Journal of Economics*, 2001, 116, 55–79.
- Chew, Soo Hong, “Axiomatic Utility Theories with the Betweenness Property,” *Annals of Operations Research*, 1989, 19, 273–298.
- Cohn, Alain, Jan Engelmann, Ernst Fehr and Michel Maréchal, “Evidence for Countercyclical Risk Aversion: An Experiment with Financial Professionals,” *UBS Center Working Paper No. 4*, 2013.
- Cohen, Alma and Liran Einav, “Estimating Risk Preferences from Deductible Choice,” *American Economic Review*, 2007, 97(3), 745–788.
- Dekel, Eddie, “An axiomatic characterization of preferences under uncertainty: Weakening the independence axiom,” *Journal of Economic Theory*, 1986, 40, 304–318.
- Gordon, Stephen and Pascal St-Amour, “A Preference Regime Model of Bull and Bear Markets,” *American Economic Review*, 2000, 90, 1019–1033.
- Guiso, Luigi, Paola Sapienza and Luigi Zingales, “Does Local Financial Development Matter?,” *Quarterly Journal of Economics*, 2004, 119, 929–969.
- Guiso, Luigi, Paola Sapienza and Luigi Zingales, “Time Varying Risk Aversion,” *mimeo*, 2011.
- Gul, Faruk, “A Theory of Disappointment Aversion,” *Econometrica*, 1991, 59, 667–686.
- Hanany, Eran, and Peter Klibanoff, “Updating Preferences with Multiple Priors,” *Theoretical Economics*, 2007, 2, 261–298.
- Kaustia, Markku and Samuli Knupfer, “Do investors overweight personal experience? Evidence from IPO subscriptions,” *The Journal of Finance*, 2008, 63(6), 2679–2702.
- Knutson, Brian and Stephanie Greer, “Anticipatory affect: Neural correlates and consequences for choice,” *Philosophical Transactions of the Royal Society B*, 2008, 363, 3771–3786.

- Köszegi, Botond, and Matthew Rabin, “Reference-Dependent Consumption Plans,” *American Economic Review*, 2009, 99, 909–936.
- Kreps, David M., and Evan L. Porteus, “Temporal Resolution of Uncertainty and Dynamic Choice Theory,” *Econometrica*, 1978, 46, 185–200.
- Kuhnen, Camelia and Brian Knutson, “The impact of anticipatory affect on beliefs, preferences, and financial risk taking,” *Journal of Financial and Quantitative Analysis*, 2011, 46, 605–626.
- Lee, Young Han and Ulrike Malmendier, “The Bidder’s Curse,” *American Economic Review*, 2011, 101, 749–787.
- Machina, Mark J., “Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty,” *Journal of Economic Literature*, 1989, 27, 1622–1668.
- Malmendier, Ulrike and Stefan Nagel, “Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?,” *Quarterly Journal of Economics*, 2011, 126(1), 373–416.
- Mas-Colell, Andreu, Michael Whinston and Jerry Green, *Microeconomic Theory*, 1995. Oxford University Press, USA.
- Parducci, Allen, *Happiness, Pleasure, and Judgment: The Contextual Theory and Its Applications*, 1995. New Jersey: Lawrence Erlbaum Associates.
- Post, Thierry, Martijn van den Assem, Guido Baltussen and Richard Thaler, “Deal or No Deal? Decision Making Under Risk in a Large-Payoff Game Show,” *American Economic Review*, 2008, 98, 38–71.
- Roese, Neal and James Olson (eds.), *What Might Have Been: The Social Psychology of Counterfactual Thinking*. Mahwah, New Jersey: Lawrence Erlbaum Associates, 1995.
- Ronen, Joshua, “Effects of Some Probability Displays on Choices,” *Organizational Behavior and Human Performance*, 1973, 9, 1–15.
- Rozen, Kareen, “Foundations of Intrinsic Habit Formation,” *Econometrica*, 2010, 78, 1341–1373.
- Saha A., C. Shumway and H. Talpaz, “Joint Estimation of Risk Preference Structure and Technology Using Expo-Power Utility” *American Journal of Agricultural Economics*, 1994, 76, 173–184.

- Schwarz, Norbert and Fritz Strack, "Reports of Subjective Well-Being: Judgmental Processes and Their Methodological Implications" in *Well-Being: The Foundations of Hedonic Psychology*, 1998, eds. D. Kahneman, E. Diener, N. Schwarz. New York: Russell Sage Foundation.
- Segal, Uzi, "Two-stage Lotteries without the Reduction Axiom," *Econometrica*, 1990, 58, 349–377.
- Shiller, Robert, "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?," *American Economic Review*, 1981, 71(3), 421–436.
- Sunstein, Cass, "Probability Neglect: Emotions, Worst Cases, and Law," *Yale Law Journal*, 2002, 112, 61–107.
- Smith, Richard, Ed Diener and Douglas Wedell, "Intrapersonal and social comparison determinants of happiness: A range-frequency analysis," *Journal of Personality and Social Psychology*, 1989, 56, 317–325.
- Thaler, Richard and Eric Johnson, "Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice," *Management Science*, 1990, 36, 643–660.
- Yesuf, Mahmud and Randall Bluffstone, "Poverty, Risk Aversion, and Path Dependence in Low-Income Countries: Experimental Evidence from Ethiopia," *Amer. J. Agr. Econ.*, 2009, 91(4), 1022-1037.