1-4-2019

A Bias Aggregation Theorem

Mark Schneider
Chapman University

A Bias Aggregation Theorem

Comments

Working Paper 19-03

A Bias Aggregation Theorem

Mark Schneider[†]

January 4, 2019

Abstract

In a market where some traders are rational (maximize expected utility) and others are systematically biased (deviate from expected utility due to some bias parameter, θ), do equilibrium prices necessarily depend on θ? In this note, focusing on the case where there is an aggregate and systematic bias in the population, we show that market prices can still be unbiased. Hence, we establish that systematically biased agents do not necessarily imply biased market prices. We show that the parametric model we use also predicts observed deviations from expected utility in laboratory and market environments.

Keywords: Risk aversion; Expected utility; Bias Aggregation
JEL Codes: D81; D90

[*]Acknowledgments: I thank Manel Baucells, Cary Deck, Enrico Diecidue, Byung-Cheol Kim, Robert Shiller, and Charlie Sprenger for comments and encouragement regarding this research. I am responsible for any errors.

[†]University of Alabama, 361 Stadium Drive, Tuscaloosa AL 35487 (e-mail: maschneider4@cba.ua.edu; phone: (205) 348-1554).

1 Introduction

The expected utility (EU) model has been the workhorse of economic analysis since it was axiomatized by John von Neumann and Oscar Morgenstern over seventy years ago.[1] It was swiftly applied to decisions involving insurance (Friedman and Savage, 1948), the existence of mixed strategy equilibria in non-cooperative games (Nash 1950a, 1951), co-operative bargaining theory (Nash, 1950b) and optimal portfolio selection (Tobin 1958; Markowitz, 1959). With the rise of information economics, EU became the micro foundation for analyzing games, markets, and contracts under asymmetric information. With the rise of modern finance, EU became the micro foundation of equilibrium asset pricing theory. It has penetrated the barriers into neighboring disciplines, finding application in fields as diverse as political science, evolutionary biology, and sociology. Even one of its greatest critics has referred to EU as "the most important theory in the social sciences" (Kahneman 2011, p. 270).

Despite its widespread adoption in economics and its normative basis as a model of rational choice, EU has been subject to persistent empirical challenges. Four of the most robust and systematic empirical violations of EU are (i) the Allais paradox, (ii) the common ratio effect (both identified by Allais, 1953), (iii) loss aversion (e.g., aversion to small mixed gambles) as defined by Kahneman and Tversky (1979), and (iv) aversion to ambiguity (e.g., as defined by Ghirardato and Marinacci, 2002).

In addition to these and other empirical limitations of EU in laboratory experiments, EU also fails to capture some features of decision making in the 'real world'. For instance, under EU, consumers would not find it optimal to purchase full insurance at actuarially unfair prices (Blavatskyy, 2014), contrary to the large premiums many people are willing to pay to eliminate risk. Sydnor (2010) found that risk aversion under EU is unable to explain individual insurance decisions for moderate risks. In portfolio choice, the EU model is unable to explain why many households do not invest in stocks, despite the historically large average return on equities. In financial markets, buying and selling price gaps in markets cannot be explained by EU (Dow and Werlang, 1992), and the standard consumption-based capital asset pricing model with EU preferences falls victim to the equity premium puzzle (Mehra and Prescott, 1985; Chateauneuf et al., 2007) and other asset pricing anomalies.

[1]See Bleichrodt et al. (2016) for the conclusion that John Nash and Jacob Marschak were the first to provide a complete axiomatization of EU in their 1950 papers in the same issue of Econometrica (Marschak 1950; Nash 1950b).

In this note, we first establish that a one-parameter extension of EU in which a decision maker exhibits 'security bias' – a disproportionate preference for lotteries with higher minimum outcomes, provides a unified explanation for the experimental evidence (the Allais paradox, common ratio effect, loss aversion, and preference for robustness to ambiguity) as well as for the empirical challenges to EU in insurance markets and portfolio choice noted above. We do so by transforming individual (rather than cumulative) probabilities. As a consequence, our approach provides a clean comparison to the standard predictions of EU.

As our main result, we embed this model of security-biased agents into a market with Arrow securities and prove a 'bias aggregation theorem' which demonstrates that a market with systematically biased agents can aggregate investor biases and produce unbiased market-level prices. That such a formal result is possible has not been clear even though it has long been argued by neoclassical economists that markets can eliminate investor biases and it has been argued by behavioral economists that markets cannot do so (Thaler, 2015).

To motivate the model of security bias, we note that despite a half-century of experimental and theoretical work on choice under risk, the standard generalization of EU that incorporates systematically biased agents has not yet been established. In a recent paper, O'Donoghue and Somerville (2018) summarize three promising approaches to modeling risk aversion in economics – loss aversion (Kahneman and Tversky, 1979; Koszegi and Rabin, 2006, 2007), probability weighting (Quiggin, 1982), and salience theory (Bordalo et al., 2012). The Koszegi-Rabin approach to loss aversion has become popular in recent years, although it makes a strong assumption that people have rational expectations regarding the reference point. It also violates stochastic dominance for a loss aversion parameter greater than 2. Salience theory is a powerful framework for analyzing behavior although it violates transitivity in general and it violates stochastic dominance when lottery outcomes are correlated (Bordalo et al., p.1259). The approach by Quiggin avoids violations of stochastic dominance, but models of rank-dependent probability weighting have not proved very tractable in economic analysis as they transform cumulative rather than individual probabilities. Yet doing so seemed necessary - as O'Donoghue and Somerville (2018) note, "While early formulations applied the probability weighting function to the probability of each outcome, this approach generates violations of dominance." (p.102).

Given the preceding discussion, one might ask whether there is an alternative approach to modeling risk aversion that avoids the limitations of the three approaches noted above. We argue that the answer to this question appears to be 'yes' and the resulting model appears to be one of 'quasi-rank dependent' utility (QRD), which can be viewed as a natural risky-choice analog to quasi-hyperbolic discounting (Laibson, 1997). The QRD model formalizes 'security bias', it satisfies basic axioms of rational choice such as transitivity and stochastic dominance, and it is simpler than standard rank-dependent probability weighting models (Quiggin, 1982) since it transforms individual rather than cumulative probabilities. This observation may be surprising since it is widely believed that models which transform individual probabilities necessarily violate stochastic dominance or transitivity (Diecidue et al., 2004; Dhami, 2016; O'Donoghue and Somerville, 2018).

In addition to accommodating observed empirical patterns, there is the related issue that some models have more empirical content than others in terms of generating precise predictions. For instance, Harless and Camerer (1994) comment, "Some theories, like EU...are too lean: They could explain the data better by allowing a few more common patterns. Other theories, such as mixed fanning and rank-dependent EU, are too fat: They allow a lot of patterns which are rarely observed." (p. 1285). The QRD model provides a compromise between EU and rank dependent utility (RDU): It allows for a few more commonly observed patterns that are not explained by EU (such as the Allais paradox and common ratio effect), while generating unambiguous comparative static predictions. The QRD model is a special case of the non-extreme outcome expected utility (NEO-EU) model (Schmidt 2000; Chateauneuf et al., 2007; Webb and Zank 2011). It is, perhaps the simplest generalization of EU that preserves transitivity, stochastic dominance, and risk aversion.

We present the QRD model in section 2, study behavioral implications in section 3, and study market implications in sections 4, 5, and 6. Proofs of propositions 1 through 5 are in the appendix. Propositions 6, 7, and 8 are proved in the main text.

2 Quasi-Rank Dependent Utility Theory

Let $X \subset \mathbb{R}$ denote a finite set of outcomes. A *lottery*, f, is a probability distribution on X. Denote the set of lotteries by $\Delta(X)$. Consider model (1) where $\underline{f}$ is the minimum outcome in the support of f, $U(f) = \sum_{x \in X} f(x)u(x)$ and $\theta \in [0,1]$:

$$V(f) = \theta U(f) + (1-\theta)u(\underline{f}). \tag{1}$$

Model (1) takes the convex combination of the expected utility and the minimum utility of the lottery and can be interpreted as a disproportionate preference for lotteries with larger minimum payoffs. We refer to this behavior as 'security bias' since lotteries with higher minimum payoffs offer greater 'security' to the decision maker by limiting the worst-case scenario. Security bias can be quantified by $1-\theta$. We show that security bias offers a unified explanation for the certainty effect (the Allais paradox and common ratio effect) and loss aversion (aversion to symmetric small mixed gambles) for choices under risk (in Section 3). If probabilities are subjective, security bias can also reflect a preference for robustness to mis-specified beliefs, and it reduces to Wald's (1950) maximin rule when $\theta = 0$, which is widely used in robust decision making under uncertainty. When precise probabilities of events are unknown, the parameter θ can be interpreted as the decision maker's degree of confidence in his beliefs. If the decision maker is fully confident in his beliefs, then $\theta = 1$, and the decision maker maximizes subjective expected utility with respect to those beliefs. If the decision maker is completely uncertain of what to believe, then $\theta = 0$, and the decision maker chooses the alternative that is most robust to different specifications of beliefs. In the intermediate case where $\theta \in (0,1)$ the decision maker cares about both maximizing EU and maximizing robustness. If probabilities are subjective, (1) satisfies the definition of ambiguity aversion from Ghirardato and Marinacci (2002).

Security bias also drives behavior in strategic situations. For instance, Chong et al. (2016) find strong support for what they call 'minimum aversion' (a tendency to avoid strategies that might yield a player her lowest possible payoff) in a large set of experimental normal form games. Other behavioral interpretations for (1) are that a decision maker overweights the worst outcome of a lottery because it is salient or because it represents the most the decision maker can be guaranteed from that lottery with certainty, or because the decision maker is pessimistic or 'expects the worse'. An axiomatic foundation for (1) is given by Webb and Zank (2011). A variant of model (1) appeared in Ellsberg (1961) which Ellsberg proposed as a simple generalization of subjective expected utilty theory (p. 664) in the same paper in which he introduced his classic paradox. A different variant of (1) was introduced in Gilboa (1988), and (1) first appeared explicitly in Schmidt (2000). An analogous model to (1) for decisions under ambiguity was characterized in Kopylov (2009). Model (1) also appears as a special case of the models in Chateauneuf et al. (2007) and Webb and Zank (2011) which allow for both optimistic and pessimistic behavior. However, despite the simplicity of (1) and its convenient properties, it has received relatively little attention in applications compared to RDU and cumulative prospect theory (CPT) due to (Tversky and Kahneman, 1992).

We note that (1) has an equivalent representation as a model of quasi-rank dependent utility. Let X_f denote the random variable induced by lottery f. Let $\succsim$ denote a preference relation on $\Delta(\mathrm{X})$, with strict preference and indifference represented by $\succ$ and $\sim$.

Definition 1 (Quasi-Rank Dependent Utility) *In Quasi-Rank Dependent (QRD) utility theory, there exists utility function, u, probability weighting function, π, with $\pi(0)=0, \pi(1)=1$, and $\sum_{x\in \mathrm{X}} \pi(f(x))=1$, and a unique parameter, $\theta \in [0,1]$, such that for any $f, g \in \Delta(\mathrm{X})$, $f \succsim g$ if and only if $V(f) \geq V(g)$, where for all x in the support of f:*

$$V(f)=\sum_{x\in \mathrm{X}} \pi(f(x))u(x), \tag{2}$$

$$\pi(f(x))=\begin{cases} 1-\theta+\theta f(x), & P(X_f < x)=0. \\ \theta f(x), & P(X_f < x)>0. \end{cases} \tag{3}$$

where u is a strictly increasing utility function and $P(X_f \leq x)$ is the cumulative distribution function of X_f.

Note that the weights sum to 1. Somewhat surprisingly, the simple probability weighting function in (3) has not yet appeared in the literature. The formula in Definition 1 is a tractable "quasi-rank-dependent" probability weighting model which satisfies both stochastic dominance and transitivity and which transforms individual rather than cumulative probabilities. In addition, it is straightforward to test expected utility theory in this setup by testing if $\theta = 1$. This approach is 'quasi-rank-dependent' since the weight is different for the probability of the lowest ranked outcome, but the probabilities of all other outcomes receive the same weight. For further intuition, consider an analogy to choice over time. Note that (2-3) is somewhat analogous to the model of quasi-hyperbolic discounting (4-5) which has discount function $d(t)$, and consumption stream $(x_0, x_1, ..., x_T)$ is evaluated as (4) where $\beta \in [0,1]$. We suggest that the quasi-rank dependent model in (2-3) may be a convenient compromise between expected utility and rank-dependent utility, analogous to how the quasi-hyperbolic model in (4-5) is a convenient compromise between discounted utility and hyperbolic discounting.

$$W(x_0, x_1, ..., x_T)=\sum_t d(t)u(x_t) \tag{4}$$

$$d(t)=\begin{cases} 1 & t=0 \\ \beta\delta^t & t>0 \end{cases}. \tag{5}$$

2.1 Certainty Preference or Security Bias?

Previous work has led to the impression that the analog to quasi-hyperbolic discounting under risk is a model of 'certainty preference' with weights assigned to probabilities depending on whether $f(x) = 1$ or $f(x) < 1$. However, such models necessarily violate either stochastic dominance or transitivity (e.g., Diecidue et al., 2004), two of the most basic rationality axioms for choice under risk, which are rarely violated in experiments (e.g., Blavatskyy (2010), Regenwetter et al. (2011) and Baillon et al. (2014)). Upon closer inspection, however, it seems such models do not capture certainty preference in an intuitive way: They allow for situations where a guaranteed payoff x is preferred to a 50-50 chance of gaining $y > x$ or x, merely because the former is certain. It seems plausible in such cases that certainty is better viewed as the minimum guarantee of a lottery: While the certain outcome guarantees x, the lottery guarantees at least x. This suggests that a more natural model of certainty preference assigns different weights to the probability of x if $P(X_f \geq x) < 1$ versus if $P(X_f \geq x) = 1$ (the lottery guarantees at least x), as implied by (2). It also appears that models which assign different weights to probabilities depending on whether $f(x) = 1$ or $f(x) < 1$ do not capture the behavior of most subjects who exhibit the Allais paradox. For instance, Incekara-Hafalir and Stecher (2016) conducted a novel test of Allais-style violations of EU in which they replaced the common consequence across a series of six decisions. Their subjects chose between safe lottery ($c, 0.89; $8, 0.10; $8, 0.01) and risky lottery ($c, 0.89; $10, 0.10; $0, 0.01) for c$\in\{0, 5, 8, 10, 16, 20\}$. Let R denote a 'risky' choice and S denote a 'safe' choice. Ordering the six choices from those with the lowest value of c to the highest value of c, only the preference patterns *RRRRRR* and *SSSSSS* are consistent with EU. The classic Allais paradox corresponds to the case where the first letter in the sequence is R (where c = 0) and the third letter is S (where c = 8). Models which assign different weights to certain and uncertain outcomes predict a 'certainty effect' pattern of *RRSRRR* but not the 'zero effect' pattern *RSSSSS*. In contrast, QRD is consistent with a 'zero effect' pattern *RSSSSS*, but not with a certainty effect pattern. The standard RDU model permits the zero effect pattern, the certainty effect pattern, and the reverse of each pattern, and so is not very helpful in predicting which effect will dominate. Consistent with the predictions of QRD, Incekara-Hafalir and Stecher (2016) observed strong support for the zero effect pattern while none of their subjects exhibited the certainty effect pattern.

3 Behavioral Implications

We next explore behavioral implications of the QRD model.

3.1 Loss Aversion

A basic property of observed behavior under risk is loss aversion which has been defined behaviorally by Kahneman and Tversky (1979) and Schmidt and Zank (2005) as aversion to 50-50 gain-loss bets: Given a choice between lotteries f and g where $f := (y, 0.5; -y, 0.5)$ and $g := (x, 0.5; -\mathrm{x}, 0.5)$, for any $x > y \geq 0$, *loss aversion* holds if $f \succ g$.

Proposition 1 *Let a decision maker have QRD preferences with $u(x) = x$. Then loss aversion holds if and only if $\theta \in [0,1)$.*

3.2 Rabin's Paradox

Rabin (2000) proved a calibration theorem which implies, for instance, that an EU maximizer who rejects a 50-50 bet to lose \$100 or win \$125 at all wealth levels will also reject a 50-50 bet to lose \$600 or win \$1 million, an implausible level of risk aversion.

Under the QRD model, the 50-50 lose \$100, gain \$125 bet is rejected at all wealth levels, $w >$ \$100, if $u(w) > (1-\theta/2)u(w-100) + (\theta/2)u(w+125)$ for all $w >$ \$100. If u is strictly concave, it follows that

$$u((1-\theta/2)(w-100) + (\theta/2)(w+125)) > (1-\theta/2)u(w-100) + (\theta/2)u(w+125)$$

for all $\theta \in (0,1)$. Also note that for any $\theta \in (0, 8/9)$ and any we have:

$$u(w) > u((1-\theta/2)(w-100) + (\theta/2)(w+125)).$$

Thus, for any strictly concave utility function, a QRD agent with $\theta \in (0, 8/9)$ will reject Rabin's small stakes gamble at any wealth level, $w >$ \$100. Next, consider the 50-50 gamble to lose \$600 or gain \$z. The gamble is accepted at current wealth level, w, if $(\theta/2)u(w+\mathrm{z}) + (1-\theta/2)u(w-600) > u(w)$. Acceptance of the gamble also implies that:

$$u(w + z(\theta/2) - 600(1-\theta/2)) > u(w),$$

which implies $\theta > 1200/(600+\mathrm{z})$. Under the preceding restrictions on θ, a QRD agent may reject the small stakes gamble at all wealth levels and choose the gamble with a 50-50 chance of losing \$600 or gaining \$z for sufficiently large z.

3.3 The Common Ratio Effect

In the following analysis, we provide general conditions on θ which explains the common ratio (at certainty), a robust violation of the independence axiom of EU.

Definition 2 *Consider lotteries* $f := (y,1), f' := (y,q;0,1-q)$, $g := (x,p;0,1-p)$, *and* $g' := (x,qp;0,1-qp)$, *for any* $y \in (0,x)$ *and* $p,q \in (0,1)$*. The common ratio effect holds if* $f \sim g$ *implies* $f' \prec g'$*.*

Proposition 2 *Let* $u(0) = 0$*. For a decision maker with QRD preferences, the common ratio effect holds if and only if* $\theta \in (0,1)$*.*

In the classic version of the common ratio effect due to Kahneman and Tversky (1979), $(x,y,p,q) =$ (\$4000, \$3000, 0.80, 0.25). Definition 2 implies that an agent who is indifferent between lotteries f and g will strictly prefer g' over f'.

3.4 The Allais Paradox

A different violation the EU independence axiom, the Allais paradox (the common consequence effect) due to Allais (1953), is defined as:

Definition 3 *Define lotteries* $f := (y,1)$, $f' := (y,q;0,1-q)$, $g := (x,p;y,1-q;0,q-p)$, *and* $g' := (x,p;0,1-p)$, *for any* $y \in (0,x)$ *and* $p,q \in (0,1)$*. The Allais paradox holds if* $f \sim g$ *implies* $f' \prec g'$*.*

Proposition 3 *Let* $u(0) = 0$*. For a decision maker with QRD preferences, the Allais paradox holds if and only if* $\theta \in (0,1)$*.*

While the Allais paradox is observed at the large stakes of Allais (1953) where $(x,y,p,q) =$ (\$5 million, \$1 million, 0.10, 0.11) and at the stakes used by Kahneman and Tversky (1979), where $(x,y,p,q) =$ (\$2500, \$2400, 0.33, 0.34), the paradox is greatly diminished at small stakes. In particular, when payoffs are scaled down to $(x,y,p,q) =$ (\$100, \$20, 0.10, 0.11) as in Fan (2002), or to $(x,y,p,q) =$ (\$25, \$5, 0.10, 0.11), as in Huck and Muller (2012), experimental subjects typically choose the riskier lottery in both choices. There is a strong intuitive basis for not observing the paradox at these small stakes: People are naturally willing to accept the 1% chance of receiving \$0 in exchange for a 10% chance of receiving \$100. A more complete explanation of the Allais paradox should predict the paradox at the large stakes observed by Allais, and Kahneman and Tversky, but should predict behavior consistent with EU at the smaller stakes used by Fan (2002) and Huck and Muller (2012). The standard version of CPT with a power value function defined over gains and losses cannot account for this aggregate pattern even given any probability weighting function. However, QRD naturally accommodates all four cases. For example, if $u(w+x) = \ln(w+x)$, for an agent with current wealth w = \$100,000, and $\theta = 0.9$, then QRD predicts the Allais paradox to be observed for the examples by Allais (1953) and Kahneman and Tversky (1979), and predicts behavior consistent with EU (the choice of the two riskier options) for the examples from Fan (2002) and Huck and Muller (2012).

4 Application to Insurance Markets

For decisions under risk, an important implication of quasi-rank dependent probability weighting is that a decision maker will pay a disproportionately higher premium for risk elimination than for an equivalent degree of risk reduction that does not eliminate the risk. This prediction also has empirical support. For instance, Botzen et al. (2013) find that households place a substantial premium on policies to eliminate flood risk relative to other opportunities which merely reduce the risk. A similar conclusion was reached by Viscusi et al. (2014) who found there to be a greater premium for policies that reduce cancer risks to zero relative to policies which reduce but do not eliminate the risk. We illustrate this preference for risk elimination in the context of insurance.

4.1 The Decision to Purchase Regular Insurance

Consider the following situation described by Blavatskyy (2018): A decision maker has a risk of losing D dollars with probability $q \in (0,1)$. The decision maker has the option of purchasing $x \in [0,D]$ units of insurance, where one unit of insurance costs c dollars, with $c \in (0,1)$, and pays the decision maker one dollar if the loss occurs. We consider the optimality of purchasing regular (full) insurance (i.e., the case where $x = D$) under the QRD model. Formally, we have the following result:

Proposition 4 *For any $\theta \in (0,1)$, and any concave u, a consumer with QRD preferences finds it optimal to purchase full insurance at an actuarially unfair price $c > q$ if (6) holds:*

$$1 - \theta + \theta q \geq c. \tag{6}$$

Under EU, $\theta = 1$ and full insurance is optimal when $q \geq c$. That is, a risk-averse EU agent will not purchase full insurance at an actuarially unfair price. In contrast, Proposition 4 implies that a QRD agent will do so if the weight the agent places on the loss, $1 - \theta + \theta q$, exceeds the cost per unit of insurance, c.

4.2 Over-insurance of Moderate Risks

The QRD model also provides a plausible explanation for the tendency to over-insure moderate risks, such as paying for highly priced extended warranties on consumer products. To illustrate overinsurance of moderate risks, consider the finding from Sydnor (2010) from data on real insurance purchases that customers with a 4 percent probability of a loss were willing to pay \$95 to lower the deductible from \$1,000 to \$500. In particular, under EU, for a constant relative risk aversion (CRRA) utility function with coefficient of relative risk aversion, r, Sydnor estimated a lower bound of $r > 1,839$ for customers who purchased the insurance contract with the \$500 deductible. This estimate is far above plausible levels of relative risk aversion (e.g., Holt and Laury, 2002). However, the majority of new customers in Sydnor's study purchased this contract.

The QRD model provides a simple explanation for Sydnor's finding. Let p denote the price of insurance with a \$1,000 deductible, and let w denote the consumer's initial wealth. For simplicity and to isolate the role of security bias, let $u(x) = x$. Then a QRD consumer prefers to pay \$95 to lower the deductible from \$1,000 to \$500 for all $\theta < 0.84375$. That is, under QRD, even a consumer with *linear* utility will purchase the contract with the \$500 deductible for a wide range of parameter values (any $\theta \in (0, 0.84375)$).

5 Application to Portfolio Choice

Consider a simple application to portfolio choice. There is a set, S, of possible states. Let q_s be the probability of state s. There is a safe asset that pays \$0 in each state and a risky asset, R, that yields return r_s in state s, with minimum return $r_{\underline{s}} < 0$. A QRD agent with wealth w, chooses an amount x to invest in the risky asset to maximize (7):

$$V(R) = \theta \sum_{s \in S} q_s u(w + xr_s) + (1 - \theta) u(w + xr_{\underline{s}}). \tag{7}$$

An EU maximizer ($\theta = 1$) will invest a positive amount in the risky asset if it has a positive expected return *regardless* of the agent's degree of risk aversion. This implication is contrary to the finding that many households have limited or no participation in the stock market (e.g., Mankiw and Zeldes, 1991). The QRD model provides an explanation for this limited stock market participation. Formally:

Proposition 5 *For any* $\theta \in (0, 1)$, *and any concave* u, *there is always sufficiently small* $\varepsilon > 0$ *such that an investor with QRD preferences will find it optimal to not invest in a risky asset,* R, *with expected return* $\mathbb{E}[R] = \varepsilon$ *and minimum return* $r_{\underline{s}} < 0$.

6 Bias Aggregation

Having shown that biases do matter in individual insurance and portfolio decisions, we now show that systematic biases do not necessarily affect market level prices. We consider a market for Arrow securities which pay \$1 if a target event occurs and \$0 otherwise. There are m traders, $j = 1,...,m$, of whom $n < m$ are buyers and $m-n$ are sellers. There is heterogeneity in beliefs, where trader j believes the target event will occur with probability q_j. Wealth, w, is assumed to be independent of beliefs. Beliefs are drawn from a distribution $\mathrm{F}(q)$. Traders are price-takers and pursue trading strategies which maximize their preferences. We assume that preferences are given by the QRD model. Traders are risk-averse with log utility. We first assume traders have the same bias parameter θ (with degree of bias given by $\hat{\theta} = 1-\theta$). We later relax this assumption to allow for heterogeneity in the degree of bias. Let p denote the price of the security and let x_j denote the quantity of the security purchased by trader j, for $j = 1,...,n$. Let y_j denote the quantity of the security sold by trader j, for $j = n+1,...,m$. Buyers of the security solve (8):

$$\max_{x_j} V_j = \theta q_j \ln(w + x_j(1-p)) + (1-\theta+\theta(1-q_j))\ln(w - x_j p) \tag{8}$$

for $j = 1,...,n$. Sellers of the security solve a similar maximization problem:

$$\max_{y_j} V_j = (1-\theta+\theta q_j)\ln(w - y_j(1-p)) + \theta(1-q_j)\ln(w + y_j p) \tag{9}$$

for $j = n+1,...,m$.The optimal quantity, x_j^*, demanded by buyer $j = 1,...,n$ and the optimal quantity, y_j^*, offered by seller $j = n+1,...,m$ are given by:

$$x_j^* = w\frac{\theta q_j - p}{p(1-p)}, \tag{10}$$

$$y_j^* = w\frac{\theta(1-q_j)-(1-p)}{p(1-p)}. \tag{11}$$

From the formulas for x_j^* and y_j^*,a buyer $j's$ demand is positive if $\theta q_j - p > 0$ and seller $j's$ supply is positive if $\theta(1-q_j)-(1-p) > 0$. In equilibrium, supply equals demand.

$$\int_{-\infty}^{1+\frac{p}{\theta}-\frac{1}{\theta}} w\frac{\theta(1-q)-(1-p)}{p(1-p)} f(q)dq = \int_{\frac{p}{\theta}}^{\infty} w\frac{\theta q - p}{p(1-p)} f(q)dq. \tag{12}$$

From (12), we see that buying-selling price gaps will exist in equilibrium.

Proposition 6 *Buying-selling price gaps exist in equilibrium: For QRD preferences, a trader with belief q will buy if and only if $p < \theta q$ and will sell if and only if $p > \theta q + 1 - \theta$.*

A similar finding was derived by Dow and Werlang (1992) in the context of portfolio choice.[2]

Note that when $\theta = 1$ (i.e., under EU preferences), there are no buying-selling price gaps: A trader with belief q will buy the security when $p < q$ and will sell when $p > q$. Also note that the size of the buying-selling price gap is equal to the degree of security bias $(1-\theta)$. For $\theta < 1$, security bias results in an efficiency loss: A buyer with belief q, could earn a positive subjective expected payoff by trading for all $q \in (p, \frac{p}{\theta})$ and a seller with belief q could earn a positive subjective expected payoff by trading for all $q \in (1 + \frac{p}{\theta} - \frac{1}{\theta}, p)$. Yet no trade takes place for agents with $q \in (1 + \frac{p}{\theta} - \frac{1}{\theta}, \frac{p}{\theta})$. In an extreme case, if security bias is sufficiently strong, it can even lead to market failure. For instance, if $p = 0.5$ and $\theta = 0.5$, then no trade occurs for any beliefs $q \in (0,1)$, even though welfare-enhancing trade could occur for all buyers with $q > p$ and all sllers with $q < p$. For a less extreme example, if $p = 0.5$ and $\theta = 0.9$, then no trade occurs for beliefs $q \in (\frac{4}{9}, \frac{5}{9})$.

If wealth is independent of beliefs, the equilibrium price under (12) is:

$$p = \frac{(1-\theta)\int_{-\infty}^{1+\frac{p}{\theta}-\frac{1}{\theta}} f(q)dq + \theta\left[\int_{-\infty}^{1+\frac{p}{\theta}-\frac{1}{\theta}} qf(q)dq + \int_{\frac{p}{\theta}}^{\infty} qf(q)dq\right]}{\int_{-\infty}^{1+\frac{p}{\theta}-\frac{1}{\theta}} f(q)dq + \int_{\frac{p}{\theta}}^{\infty} f(q)dq}.$$

To consider the impact of security bias on market prices, let $\overline{q}$ denote the average belief across the entire population of traders[3] (i.e., $\overline{q} = \int_{-\infty}^{\infty} qf(q)dq$).

Definition 4 *The equilibrium market price, p, is unbiased if $p = \overline{q}$.*

Gjerstand (2005) and Wolfers and Zitzewitz (2006) show that in the EU case (i.e., when $\theta = 1$), market prices are unbiased. A natural question is whether prices can be unbiased

[2]Dow and Werlang (1992) demonstrated that buying-selling price gaps for assets in financial markets hold under the more general Choquet Expected Utility preferences. The model of Bordalo et al. (2012a) generates buying-selling price gaps in the context of salience-based consumer choice.

[3]One could alternatively define $\overline{q}$ as the average belief across all traders whose demand or supply functions are positive. The conclusions in Propositions 7 and 8 continue to hold under this alternative specification of $\overline{q}$.

despite the systematic security bias of market participants (for $\theta < 1$). That is, do equilibrium prices necessarily depend on θ if traders have preferences that depend on θ? We offer a simple bias aggregation theorem:

Proposition 7 *Suppose buyers and sellers maximize* (8) *and* (9), *respectively, and that supply and demand are each positive. If beliefs across agents are uniformly distributed on* $[0,1]$, *then equilibrium market level prices wil be unbiased* $(p = \overline{q})$ *for any systematic individual level bias* $\theta \in (0.5, 1]$.

The restriction $\theta \in (0.5, 1]$ is necessary for supply and demand to each be positive.[4] Given our assumption that wealth is independent of beliefs, if $f(q)$ is uniform on [0,1], and supply and demand are positive, then (12) becomes:

$$\frac{w}{p(1-p)} \int_0^{1+\frac{p}{\theta}-\frac{1}{\theta}} (\theta(1-q) - (1-p))dq = \frac{w}{p(1-p)} \int_{\frac{p}{\theta}}^{1} (\theta q - p)dq.$$

The equilibrium price is the price p that solves $(p+\theta-1)^2 = (p-\theta)^2$, yielding:

$$p = \frac{1-2\theta}{2-4\theta} = \frac{1}{2} = \overline{q}.$$

Although simple, Proposition 7 is a surprising result - aggregating preferences of *systematically biased* agents produces *unbiased* market level prices!

A primary finding in behavioral economics is that biases are systematic. Thus, they will not cancel out as noise in ways that random errors might. In contrast, many economists argue that biases will be eliminated by the market. In his book, Misbehaving: The Making of Behavioral Economics, Thaler (2015) writes: "I call this argument the invisible handwave. . . The vague argument is that markets somehow discipline people who are misbehaving. Handwaving is a must because there is no logical way to arrive at a conclusion that markets transform people into rational agents." Interestingly, Proposition 7 considers the case where biases are systematic in the same direction (in a manner consistent with the Allais paradox, and loss aversion) and shows that even if all agents in a market are systematically biased, equilibrium prices can accurately aggregate beliefs and produce the same prices as-if all agents maximized expected utility. Proposition 7 does not mean that biases do not affect market prices in general. The result is restrictive, particularly given the

[4] Given that beliefs are uniformly distributed over [0,1], the restriction that supply and demand are positive implies that $[1+\frac{p}{\theta}-\frac{1}{\theta}, \frac{p}{\theta}] \subset [0,1]$ or, equivalently, that $1-\theta < p < \theta$.

uniform assumption for beliefs and the assumption that traders have the same risk aversion and bias (although one might view this as a representative agent).

To push Proposition 7 a little further, suppose that in addition to heterogeneity in subjective beliefs, q_j, there is also heterogeneity in the bias parameter, θ_j. Let $\hat{\theta}_j := 1-\theta_j$ denote the degree of security bias for agent j. Hence, an unbiased agent has $\hat{\theta}_j = 0$ and so maximizes expected utility. To allow for heterogeneity in security bias, suppose that θ_j is uniformly distributed over $[k,1]$,where $k \in (0.5,1)$. The restriction that $k > 0.5$ ensures that the supply and demand functions are positive for all agents. Note that this heterogeneity in biases does *not* average over agents who have security bias and agents who have the 'reverse' bias. All agents still have systematic security bias (e.g., they all exhibit the Allais paradox and the common ratio effect, except for agents with $\hat{\theta}_j = 0$ who exhibit no bias). Hence, the 'average' bias in the market is some value, $\mathbb{E}[\hat{\theta}] > 0$.

Proposition 8 *Let $k \in (0.5,1)$ and suppose buyers and sellers maximize* (8) *and* (9), *respectively. If beliefs, q_j are uniformly distributed on* $[0,1]$, *and security bias $\hat{\theta}_j$ across agents is uniformly distributed on* $[0,1-k]$, *then equilibrium market level prices will be unbiased* ($p=\overline{q}$) *even though all agents are systematically biased, with average bias* $\mathbb{E}[\hat{\theta}] > 0$.

Allowing for heterogeneity in both subjective probabilities and bias parameters, with q_j uniformly distributed over [0,1] and θ_j uniformly distributed over $[k,1]$, the market clearing condition reduces to:

$$\int_k^1 \int_0^{1+\frac{p}{\theta}-\frac{1}{\theta}} (\theta(1-q)-(1-p))dqd\theta = \int_k^1 \int_{\frac{p}{\theta}}^1 (\theta q - p)dqd\theta.$$

After integrating over subjective probabilities, the above condition becomes:

$$\int_k^1 \frac{(\theta+p-1)^2}{2\theta}d\theta = \int_k^1 \frac{(\theta-p)^2}{2\theta}d\theta. \tag{13}$$

This condition simplifies to:

$$-\frac{(p-1)^2 \ln k}{2} + (p-1)(1-k) = -\frac{p^2 \ln k}{2} - p(1-k),$$

which further simplifies to $p = \frac{1}{2} = \overline{q}$.

A subtle point related to Proposition 7 is that traders may revise their prior beliefs given the information revealed by market prices to form posterior beliefs. We have implicitly assumed that traders have fixed beliefs (i.e., their prior and posterior probabilities are equal).

We next provide a simple and plausible illustration that equilibrium prices can be the same for prior and posterior beliefs, even if traders update their prior beliefs based on the information they extract from observing market prices. To distinguish prior and posterior beliefs, we now denote the former by q and the latter by $q(p)$. Consider a simple and plausible belief-updating rule from Manski (2006) in which a trader's posterior belief is a weighted average of her prior belief and the observed market price. That is:

$$q(p) = \lambda q + (1-\lambda)p, \tag{14}$$

where $\lambda \in [0,1]$. Under this rule, prior beliefs are updated in the direction of the market price, with $(1-\lambda)$ determining the degree to which beliefs are revised, including as special cases no revision ($\lambda = 1$) and full revision ($\lambda = 0$). If prior beliefs are uniformly distributed over $[0,1]$,then the distribution of posterior beliefs is uniform over the interval $[(1-\lambda)p, \lambda + (1-\lambda)p]$ which is a subset of $[0,1]$. Performing the same analysis as in Proposition 7, assuming supply and demand are each positive, in equilibrium, the aggregate quantity supplied is given by (15):

$$\frac{w}{p(1-p)} \int_{(1-\lambda)p}^{\frac{p(1-\theta(1-\lambda))-(1-\theta)}{\theta\lambda}} \frac{(\theta(1-p-\lambda(q-p))-(1-p))}{\lambda} dq \tag{15}$$

The aggregate quantity demanded is given by (16):

$$\frac{w}{p(1-p)} \int_{\frac{p(1-\theta(1-\lambda))}{\theta\lambda}}^{\lambda+(1-\lambda)p} \frac{(\theta(p+\lambda(q-p))-p)}{\lambda} dq \tag{16}$$

From setting supply equal to demand in equilibrium, we find that the equilibrium price is the price p that solves the following equation:

$$(p(\lambda^2-1)\theta + p + (\theta-1))^2 = (p(\lambda^2-1)\theta + p - \lambda^2\theta)^2. \tag{17}$$

Solving for p yields:

$$p = \frac{1-2\theta-\lambda^4\theta^2+\theta^2}{2(1-2\theta-\lambda^4\theta^2+\theta^2)} = \frac{1}{2} = \frac{2(1-\lambda)p+\lambda}{2} = \bar{q}. \tag{18}$$

Thus, even if agents use an updating rule that takes a weighted average of their prior and the market price to form their posterior beliefs, the equilibrium price is the same as in Proposition 7.

7 Security Bias and Expectations-based Loss Aversion

One alternative model that has gained popularity in applications is the Koszegi-Rabin model of reference-dependent preferences (Koszegi and Rabin, 2006; 2007). That model can also explain loss aversion and the Allais paradox. The loss-aversion parameter, in the Koszegi-Rabin model is consistent with rank dependent utility theory for values less than 2, but can violate stochastic dominance for values greater than 2. The model also makes the strong assumption that agents have rational expectations regarding the reference point. In contrast, QRD does not violate stochastic dominance and makes no assumptions regarding rational expectations. Under QRD there is a natural reference point that is directly observable and well-specified for each lottery – the minimum outcome in the lottery's support.

In addition to the results established for QRD in this paper, it has been shown that, QRD can generate both a large equity premium and a low risk-free rate (Chateauneuf et al., 2007; Zimper, 2012). In contrast, when the Koszegi-Rabin model is applied to generate a large equity premium, it generates 'counterfactually high volatility in the risk-free rate' (Pagel, 2016). As a consequence, a different reference point must be chosen for prospect theory applications in finance (e.g., Barberis et al., 2001), than for prospect theory applications in other domains that employ the Koszegi-Rabin model.

8 Conclusion

We motivated quasi-rank dependent (QRD) utility theory as a modeling tool for studying agents in markets in the presence of risk. The QRD model satisfies transitivity, stochastic dominance, and risk aversion, and is simpler than leading probability weighting models since it transforms individual rather than cumulative probabilities. In addition to explaining experimental paradoxes for EU, we provided results that resolve several limitations of EU when applied to real market contexts. In particular, the QRD model can simultaneously explain four of the major empirical findings in applications that cannot be explained by EU: (i) the optimality of purchasing full insurance at actuarially unfair prices; (ii) the limited stock market participation puzzle; (iii) a large equity premium (Chateauneuf et al., 2007); and (iv) the existence of buying and selling price gaps in markets (Dow and Werlang, 1992).

After demonstrating how the QRD model can explain (i) and (ii), we provided a bias aggregation theorem which shows that there are cases where market prices may be unbiased even if all traders are systematically biased. As the preceding results indicate, the QRD model provides a useful tool for analyzing markets with systematically biased agents.

Appendix: Proofs

Proposition 1. *Let $u(x) = x$. Then loss aversion holds if and only if $\theta \in [0, 1)$.*
Proof: Under the QRD model, loss aversion holds if and only if $(1 - \theta/2) > (\theta/2)$ which holds if and only if $\theta \in [0, 1)$. $\square$

Proposition 2. *The common ratio effect holds if and only if $\theta \in (0, 1)$.*
Proof: Note that for Definition 2, the QRD model implies:
$f \sim g \iff u(y) = \theta(pu(x) + (1 - p)u(0))$.
Given our normalization, $u(0) = 0$, that we noted at the beginning of Section 3, we have $u(y) > 0$ since $y > 0$ and u is strictly increasing. Hence, the above indifference can hold only if $\theta > 0$. Also,
$f' \prec g' \iff \theta(qu(y) + (1 - q)u(0)) < \theta(qpu(x) + (1 - qp)u(0))$.
Note $f' \prec g' \iff \theta u(y) < \theta(pu(x) + (1 - p)u(0)) = u(y)$.
This condition holds under QRD if and only if $\theta \in (0, 1)$. $\square$

Proposition 3. *The Allais paradox holds if and only if $\theta \in (0, 1)$.*
Proof: Note that for Definition 3, the QRD model implies:
$f \sim g \iff u(y) = \theta(pu(x) + (1 - q)u(y) + (q - p)u(0))$.
$f' \prec g' \iff \theta(qu(y) + (1 - q)u(0)) < \theta(pu(x) + (1 - p)u(0))$.
Note $f' \prec g' \iff \theta qu(y) < \theta(pu(x) + (q - p)u(0))$.
Adding $(1 - q)\theta u(y)$ to both sides yields:
$\theta u(y) < \theta(pu(x) + (1 - q)u(y) + (q - p)u(0)) = u(y)$.
Since $u(y) > 0$, this inequality holds if and only if $\theta \in (0, 1)$. $\square$

Proposition 4. *For any $\theta \in (0, 1)$, and any concave u, a consumer with QRD preferences will find it optimal to purchase regular insurance at an actuarially unfair price $c > q$ if* (6) *holds:*

$$1 - \theta + \theta q \geq c. \tag{6}$$

Proof: As noted by Blavatskyy, under regular insurance, the decision maker loses exactly cD dollars regardless of whether the loss occurs. If $x < D$, the decision maker loses cx dollars with probability $1 - q$ and loses $D + x(c - 1)$dollars with probability q. Under the quasi-rank dependent model, the decision maker will purchase regular insurance if and only if the following inequality holds for any $x \in [0, D)$.

$$u(-cD) > \theta(1-q)u(-cx) + (1-\theta+\theta q)u(-D-x(c-1)).$$

The above inequality can be arranged as follows:

$$\frac{u(-cD) - u(-D-x(c-1))}{(D-x)(1-c)}(1-c) > \frac{u(-cx) - u(-D-x(c-1))}{D-x}(1-q)\theta.$$

The fraction on the left-hand side of the above inequality is the slope of the utility function between points $-cD$ and $-D-x(c-1)$, and the fraction on the right-hand side is the slope of the utility function between points $-cx$ and $-D-x(c-1)$. For any strictly concave utility function u, the slope on the left-hand side is always greater than the slope on the right hand side of the inequality. Therefore, regular insurance will be optimal to purchase when $(1-c)/(1-q) \geq \theta$. This inequality can be rewritten as (6). □

Proposition 5. *For any $\theta \in (0, 1)$, and any concave u, there is always sufficiently small $\epsilon > 0$ such that an investor with QRD preferences will find it optimal to not invest in a risky asset, R, with expected return $\mathbb{E}[R] = \epsilon$ and minimum return $r_{\underline{s}} < 0$.*

Proof: An agent with QRD preferences and wealth w, chooses an amount x to invest in the asset in order to maximize the following value function:

$$V(R) = \theta \sum_{s \in S} q_s u(w + xr_s) + (1-\theta)u(w + xr_{\underline{s}}).$$

The first derivative of $V(R)$ with respect to x is:

$$V'(R) = \theta \sum_{s \in S} q_s u'(w + xr_s)r_s + (1-\theta)u'(w + xr_{\underline{s}})r_{\underline{s}}.$$

Consider the change in $V(R)$ from investing $x = 0$ to investing $x^* > 0$ in asset R. In that case, computing the derivative at $x = 0$ yields:

$$u'(w)\left[\theta\sum_{s\in S} q_s r_s + (1-\theta)r_{\underline{s}}\right]$$

Let $\mathbb{E}[R]$ denote the expected return on asset R. If $\theta = 1$, the above expression becomes $u'(w)\mathbb{E}[R]$. If $\mathbb{E}[R] > 0$, then the above inequality implies the classical result that an expected utility maximizer will always invest a positive amount in an asset with a positive expected return *regardless* of the agent's degree of risk aversion. Next, let $\theta < 1$ and let $\mathbb{E}[R] = 0$. Then it is clear that $\theta\sum_{s\in S} q_s r_s + (1-\theta)r_{\underline{s}} < 0$ as more weight is shifted to the lowest return. It follows that for assets with sufficiently small but positive expected returns, a QRD agent will choose not to invest in the asset. $\square$

References

1. Allais, M. 1953. "Le comportement de l'homme rationnel devant le risque: Critiques des postulats et axioms de l'ecole americaine." *Econometrica*, 21, 503-546.

2. Baillon, A., H. Bleichrodt, A. Cillo. 2014. "A tailor-made test of intransitive choice." *Operations Research*, 63, 198-211.

3. Barberis, N., M. Huang, T. Santos. 2001. "Prospect theory and asset prices." *Quarterly Journal of Economics*, 116, 1-53.

4. Blavatskyy, P. 2010. "Modifying the mean-variance approach to avoid violations of stochastic dominance." *Management Science,* 56, 2050-2057.

5. Blavatskyy, P. 2018. "A second-generation disappointment aversion theory of decision making under risk." *Theory and Decision*, 84(1): 29-60.

6. Bleichrodt, H., U. Schmidt. 2002. "A context-dependent model of the gambling effect." *Management Science*, 48, 802-812.

7. Bleichrodt, H., L. Chen, I. Moscati, P.P. Wakker. 2016. "Nash was a first to axiomatize expected utility." *Theory and Decision*, forthcoming.

8. Bordalo, P., N. Gennaioli, A. Shleifer. 2012a. "Salience and experimental tests of the endowment effect." *American Economic Review*, 102, 47-52.

9. Bordalo, P., N. Gennaioli, A. Shleifer. 2012b. "Salience theory of choice under risk." *Quarterly Journal of Economics*, 127, 1243-1285.

10. Botzen, W.J., J.C. Aerts, J.C. van den Bergh. 2013. "Individual preferences for reducing flood risk to near zero through elevation." *Mitigation and Adaptation Strategies for Global Change*. 18, 229-244.

11. Chateauneuf, A., J. Eichberger, S. Grant. 2007. "Choice under uncertainty with the best and worst in mind: Neo-additive capacities." *Journal of Economic Theory*, 137, 538-567.

12. Chong, J-K., T-H. Ho, C. Camerer. 2016. "A generalized cognitive hierarchy model of games," *Games and Economic Behavior*, 99, 257-274.

13. Dhami, S. 2016. *Foundations of Behavioral Economic Analysis*. Oxford University Press.

14. Diecidue, E., U. Schmidt, P.P. Wakker. 2004. "The utility of gambling reconsidered," *Journal of Risk and Uncertainty*, 29, 241-259.

15. Dow, J., S.R. Werlang. 1992. "Uncertainty aversion, risk aversion, and the optimal choice of portfolio." *Econometrica*, 60, 197-204.

16. Ellsberg, D. 1961. "Risk, ambiguity, and the Savage axioms." *Quarterly Journal of Economics*, 75, 643-669.

17. Fan, C-P. 2002. "Allais paradox in the small." *Journal of Economic Behavior and Organization*, 49, 411-421.

18. Friedman, M., L.J. Savage. 1948. "The utility analysis of choices involving risk." *Journal of Political Economy*, 56, 279-304.

19. Ghirardato, P., M. Marinacci. 2002. "Ambiguity made precise: A comparative foundation." *Journal of Economic Theory*, 102, 251-289.

20. Gilboa, I. 1988. "A combination of expected utility and maxmin decision criteria." *Journal of Mathematical Psychology*, 32, 405-420.

21. Gjerstad, S. 2005. "Risk aversion, beliefs, and prediction market equilibrium." Manuscript.

22. Harless, D.W., C.F. Camerer. 1994. "The predictive utility of generalized expected utility theories." *Econometrica*, 62, 1251-1289.

23. Huck, S., W. Muller. 2012. "Allais for all: Revisiting the paradox in a large representative sample." *Journal of Risk and Uncertainty*, 44, 261-293.

24. Incekara-Hafalir, E., J.D. Stecher. 2016. "A zero effect explains deviations from expected utility in Allais-type tasks." Manuscript, Carnegie Mellon University.

25. Kahneman, D. 2011. *Thinking, Fast and Slow.* New York. Farrar, Straus, Giroux.

26. Kahneman, D., A. Tversky. 1979. "Prospect theory: An analysis of decision under risk." *Econometrica*, 47, 263-292.

27. Koszegi, B., M. Rabin. 2006. "A model of reference-dependent preferences." *Quarterly Journal of Economics*, 121, 1133-1165.

28. Koszegi, B. M. Rabin. 2007. "Reference-dependent risk attitudes." *American Economic Review*, 97, 1047-1073.

29. Kopylov, I. 2009. "Choice deferral and ambiguity aversion." *Theoretical Economics*, 4, 199-225.

30. Laibson, D. 1997: "Golden eggs and hyperbolic discounting," *Quarterly Journal of Economics*, 112, 443-477.

31. Mankiw, N., S. Zeldes. 1991. "The consumption of stockholders and nonstockholders." *Journal of Financial Economics*, 29, 97-112.

32. Manski, C.F. 2006. "Interpreting the predictions of prediction markets." *Economics Letters*, 91, 425-429.

33. Markowitz, H.M. 1959. *Portfolio selection: Efficient diversification of Investments.* Yale University Press.

34. Marschak, J. 1950. "Rational behavior, uncertain prospects, and measurable utility." *Econometrica*, 18, 111-141.

35. Mehra, R., E.C. Prescott. 1985. "The equity premium: A puzzle." *Journal of Monetary Economics*, 15, 145-161.

36. Nash, J.F. 1950a. "Equilibrium points in n-person games." *Proceedings of the National Academy of Sciences*, USA, 36, 48-49.

37. Nash, J.F. 1950b. "The bargaining problem." *Econometrica*, 18, 155-162.

38. Nash, J.F. 1951. "Non-cooperative games." *Annals of Mathematics*, 54, 286-295.

39. O'Donoghue, T., J. Somerville. 2018. "Modeling risk aversion in economics." *Journal of Economic Perspectives*, 32, 91-114.

40. Pagel, M. 2016. "Expectations-based reference dependent preferences and asset pricing." *Journal of the European Economic Association*, 14, 468-514.

41. Quiggin, J. 1982. "A theory of anticipated utility." *Journal of Economic Behavior & Organization*, 3, 323-343.

42. Regenwetter, M., J. Dana, C.P. Davis-Stober. 2011. "Transitivity of preferences." *Psychological Review*, 118, 42-56.

43. Schmidt, U. 1998. "A measurement of the certainty effect," *Journal of Mathematical Psychology*, 42, 32-47.

44. Schmidt, U. 2000. "The certainty effect and boundary effects with transformed probabilities." *Economics Letters*, 29-33.

45. Schmidt, U., H. Zank. 2005. "What is Loss Aversion?" *Journal of Risk and Uncertainty*, 30, 157-167.

46. Schneider, M. 2018. "A Capital Asset Pricing Model with Investor Sentiment." Manuscript.

47. Sydnor, J. 2010. "(Over)insuring modest risks." *American Economic Journal: Applied Economics*, 2, 177-199.

48. Thaler, R. 2015. *Misbehaving: The Making of Behavioral Economics.* Norton.

49. Tobin, J. 1958. "Liquidity preference as behavior towards risk." *Review of Economic Studies*, 25, 65-86.

50. Tversky, A., D. Kahneman. 1992. "Advances in Prospect theory: Cumulative representation of uncertainty." *Journal of Risk and Uncertainty*, 5, 297-323.

51. Viscusi, W.K., J. Huber, J. Bell. 2014. "Assessing whether there is a cancer premium for the value of a statistical life." *Health Economics*, 23, 384-396.

52. Wald, A. 1950. *Statistical decision functions.* Wiley, New York.

53. Webb, C.S., H. Zank. 2011. "Accounting for optimism and pessimism in expected utility." *Journal of Mathematical Economics*, 47, 706-717.

54. Wolfers, J., E. Zitzewitz. 2006. "Interpreting prediction market prices as probabilities." *IZA Discussion Papers*, No. 2092.

55. Zimper, A. 2012. "Asset pricing in a Lucas fruit-tree economy with the best and worst in mind," *Journal of Economic Dynamics and Control*, 36, 610-628.