Beyond Expected Utility: Decision Theory at the Frontier of Financial Modeling*

Agustín Shehadi Candela

QUAFI Research

research@quafi.io

June 2026

Abstract

Every asset price is, through the no-arbitrage Euler equation, a statement about *preferences*: the price of a payoff is its expectation under a stochastic discount factor built from the marginal utility of the investor who holds it. This survey maps the decision-theoretic frontier of that statement — how the field has moved beyond the von Neumann–Morgenstern and Savage expected-utility core in response to the systematic ways real markets and real investors *break* it. We organize the frontier around four breaks. (i) *Risk attitudes*: the Allais paradox falsifies the independence axiom, motivating rank-dependent (Quiggin, 1982; Yaari, 1987), prospect-theoretic (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992), disappointment-averse (Gul, 1991), and reference-dependent (Kőszegi and Rabin, 2006) utilities, with asset-pricing payoffs in Barberis et al. (2001); Barberis and Huang (2008). (ii) *Ambiguity*: the Ellsberg paradox falsifies a unique subjective prior, motivating Choquet (Schmeidler, 1989), maxmin (Gilboa and Schmeidler, 1989), variational (Maccheroni et al., 2006), and smooth (Klibanoff et al., 2005) preferences and robust control (Hansen and Sargent, 2001). (iii) *Time and customs*: discounted utility fails dynamic consistency (Strotz, 1955; Laibson, 1997), motivating recursive utility (Kreps and Porteus, 1978; Epstein and Zin, 1989) and habit formation (Constantinides, 1990; Abel, 1990; Campbell and Cochrane, 1999) — the formalization of "custom" as a moving reference standard of consumption. (iv) *Frictions*: liquidity (Amihud and Mendelson, 1986; Acharya and Pedersen, 2005) and asymmetric information (Akerlof, 1970; Grossman and Stiglitz, 1980; Kyle, 1985) break the representative, fully-informed agent, dissolving the very notion of a single marginal utility. Throughout, the unifying object is the pricing kernel: each behavioral feature or friction is

*QUAFI Working Paper Series, No. 2026-04.

A self-contained, rigorous survey of the decision-theoretic frontier in financial modeling. It is preliminary, circulated to inform and to invite comment; the views are the author's own. *Nothing herein is investment advice.* All representation theorems are attributed to their original published sources and stated without proof; elementary derivations are given. Any errors are mine.

a precise modification of it, and each is judged by the asset-pricing puzzles it does or does not resolve — above all the equity-premium puzzle (Mehra and Prescott, 1985) and the Hansen–Jagannathan volatility bound (Hansen and Jagannathan, 1991). We close with the open problem that ties the frontier together: the identification of preferences from data.

Keywords: decision theory; expected utility; prospect theory; ambiguity aversion; recursive utility; habit formation; liquidity; asymmetric information; stochastic discount factor; equity premium puzzle.
JEL: D81, D82, D91, G11, G12, G40.

Contents

1 Introduction: prices are statements about preferences

Asset pricing has a single organizing identity. Under the mild requirement of no arbitrage, there exists a positive random variable m_{t+1} — the *stochastic discount factor* (SDF), or pricing kernel — such that the price p_t of any traded payoff x_{t+1} satisfies

$$p_t = \mathbb{E}_t[m_{t+1}\, x_{t+1}]\,, \qquad \text{equivalently} \qquad 1 = \mathbb{E}_t[m_{t+1}\, R_{t+1}] \tag{1}$$

for gross returns $R_{t+1} = x_{t+1}/p_t$ (Cochrane, 2005). In a model with an optimizing investor, the SDF is the investor's *intertemporal marginal rate of substitution*, $m_{t+1} = \beta\, u'(c_{t+1})/u'(c_t)$. Equation (1) is therefore not a financial fact layered on top of preferences; it *is* preferences, read off prices. To model asset prices is to model the utility function inside m.

This survey is a map of how that utility function has been generalized. The classical core — expected utility (von Neumann and Morgenstern, 1944) and its subjective-probability extension (Savage, 1954) — is elegant, axiomatic, and *empirically broken* in four distinct, well-documented ways. Each break has generated a precise decision-theoretic technology to repair it, and each technology a corresponding modification of the pricing kernel. We organize the frontier accordingly: Section 2 states the expected-utility core and the puzzles that falsify its simplest form; Sections 3–6 treat the four breaks — risk attitudes, ambiguity, time & customs, and frictions (liquidity and asymmetric information); Section 7 synthesizes them as a table from break to preference technology to priced consequence; and Section 8 states the open frontier, centered on the identification problem. The aim is to be concise and pedagogic while remaining exact: every nonstandard object is defined, and every representation is attributed.

2 The expected-utility core and the puzzles that break it

2.1 Axioms and the von Neumann–Morgenstern representation

Let $\succsim$ be a preference relation over lotteries (probability measures) on a set of outcomes $X \subseteq \mathbb{R}$. The classical theory rests on four axioms.

Axiom 1 (The expected-utility axioms)**.** $\succsim$ is *complete* and *transitive*; *continuous* (the sets $\{q : q \succsim p\}$ and $\{q : p \succsim q\}$ are closed); and satisfies *independence*: for all lotteries p, q, r and $\alpha \in (0, 1]$, $p \succsim q \iff \alpha p + (1-\alpha) r \succsim \alpha q + (1-\alpha) r$.

Theorem 1 (von Neumann–Morgenstern)**.** *$\succsim$ satisfies Axiom 1 if and only if there is a utility index $u : X \to \mathbb{R}$, unique up to a positive affine transformation, such that $p \succsim q \iff \mathbb{E}_p[u] \geq \mathbb{E}_q[u]$ (von Neumann and Morgenstern, 1944). Replacing objective lotteries by acts over states and adding the sure-thing principle yields a* subjective *probability $\mathbb{P}$ and the criterion $\mathbb{E}_{\mathbb{P}}[u]$ (Savage, 1954).*

Independence is the load-bearing axiom: it makes utility linear in probabilities and is exactly what the behavioral evidence will dispute.

2.2 Risk aversion

For a concave u, the *Arrow–Pratt* coefficient of absolute risk aversion is $A(w) = -u''(w)/u'(w)$, and of relative risk aversion $\gamma(w) = wA(w)$; the risk premium an agent pays to avoid a small gamble is proportional to $A(w)$ (Pratt, 1964). The constant-relative-risk-aversion (CRRA) family $u(c) = c^{1-\gamma}/(1-\gamma)$ is the workhorse, with γ governing both risk aversion and (its reciprocal) the elasticity of intertemporal substitution — a conflation Section 5 shows to be costly.

2.3 The pricing kernel and the equity-premium puzzle

With CRRA utility, the SDF is $m_{t+1} = \beta\,(c_{t+1}/c_t)^{-\gamma}$. Taking unconditional expectations of (1) for an excess return R^e (so $\mathbb{E}[mR^e] = 0$) and applying Cauchy–Schwarz gives the *Hansen–Jagannathan bound* (Hansen and Jagannathan, 1991):

$$\frac{|\mathbb{E}[R^e]|}{\sigma(R^e)} \leq \frac{\sigma(m)}{\mathbb{E}[m]}. \tag{2}$$

The left side is the maximal Sharpe ratio; the right side is the SDF's coefficient of variation. Empirically the U.S. equity Sharpe ratio is ≈ 0.4, while consumption growth is smooth ($\sigma(\Delta \log c) \approx$ 1–2%), so under CRRA the bound (2) requires a risk-aversion coefficient γ of order 50 or more. This is the *equity-premium puzzle* (Mehra and Prescott, 1985); forcing γ that high then implies a counterfactually high risk-free rate — the *risk-free-rate puzzle* (Weil, 1989). The canonical, frictionless, expected-utility representative agent cannot match the data with plausible preferences. Everything that follows is a disciplined response.

Theorem 2 (Rabin's calibration paradox)**.** *Rabin (2000) shows that within expected utility over* wealth*, the modest risk aversion observed over small stakes implies absurd risk aversion over large stakes: an agent who turns down a 50–50 gain-\$110 / lose-\$100 gamble at all wealth levels must reject* any *50–50 gamble that loses \$1,000, however large the possible gain. Loss aversion over* changes *(Section 3), not curvature over levels, is the natural escape.*

3 Break I — risk attitudes beyond independence (the Allais break)

3.1 The Allais paradox

Allais (1953) exhibited choices — the "common-consequence" and "common-ratio" effects — that systematically violate independence: people overweight certainty and treat a shared component of two lotteries non-neutrally, contrary to Axiom 1. Independence is descriptively false. The non-expected utility program keeps completeness, transitivity, and continuity but *weakens* independence.

3.2 Rank-dependent and dual utility

The first repair distorts cumulative *probabilities* rather than outcomes.

Definition 1 (Rank-dependent / anticipated utility)**.** With outcomes ordered $x_1 \leq \cdots \leq x_n$ and a strictly increasing weighting function $w : [0,1] \to [0,1]$, $w(0) = 0$, $w(1) = 1$, the value of a lottery is $\sum_i \pi_i \, u(x_i)$ with decision weights $\pi_i = w\big(\sum_{j \geq i} p_j\big) - w\big(\sum_{j > i} p_j\big)$ (Quiggin, 1982). Yaari (1987)'s *dual theory* is the special case $u(x) = x$, where all departure from risk-neutrality lives in w.

An inverse-S weighting w — overweighting the tails — generates the simultaneous demand for lottery-like (positively skewed) and insurance-like payoffs that linear-probability EU cannot. Machina (1982) showed, more generally, that much of EU's analytic apparatus survives if utility is merely *smooth* (Fréchet-differentiable) in the lottery, without independence.

3.3 Prospect theory

The dominant behavioral model combines probability weighting with *reference dependence* and *loss aversion.*

Definition 2 (Cumulative prospect theory)**.** Outcomes are gains/losses relative to a reference point. The value function

$$v(x) = \begin{cases} x^{\alpha}, & x \geq 0, \\ -\lambda(-x)^{\beta}, & x < 0, \end{cases} \qquad 0 < \alpha, \beta \leq 1, \;\; \lambda > 1,$$

is concave over gains, convex over losses, and *kinked* at 0 with loss-aversion coefficient λ; gains and losses are aggregated with rank-dependent weights from separate weighting functions w^{+}, w^{-} (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992).

The kink ($\lambda > 1$) is first-order risk aversion — aversion that does not vanish for small gambles — which resolves Rabin's paradox (Theorem 2) and, applied to portfolio choice, the equity premium. Kőszegi and Rabin (2006) endogenize the reference point as rational expectations of outcomes, closing the model's main degree of freedom. Gul (1991)'s *disappointment aversion* delivers first-order risk aversion within a single-parameter, dynamically tractable generalization of EU, and Loomes and Sugden (1982)'s *regret theory* derives apparent non-EU behavior from comparisons to foregone alternatives.

3.4 Asset-pricing consequences

Benartzi and Thaler (1995) show that an investor who is loss-averse and evaluates her portfolio frequently (*myopic loss aversion*) demands the observed equity premium. Barberis et al. (2001) embed prospect theory — with a reference point and changing loss aversion after gains/losses (the "house-money" effect) — in a general-equilibrium SDF, generating high, volatile, and predictable returns from smooth consumption. Barberis and Huang (2008) show that probability weighting makes positively skewed "lottery" stocks *overpriced*, earning low average returns — a preference-based account of the idiosyncratic-volatility and skewness anomalies.

4 Break II — ambiguity and model uncertainty (the Ellsberg break)

4.1 The Ellsberg paradox

Ellsberg (1961) showed that people prefer bets on *known* probabilities to bets on *unknown* ones, even when this is inconsistent with *any* single subjective prior — violating Savage's sure-thing principle. This is the formal content of Knight's distinction between *risk* (a known distribution) and *uncertainty* or *ambiguity* (an unknown one). The repair replaces the single prior of Theorem 1 with a *set* of priors.

Theorem 3 (Maxmin expected utility)**.** *Under axioms weakening Savage's (notably replacing independence with* certainty-independence *and adding* uncertainty aversion*), preferences are represented by*

$$V(f) = \min_{\mathbb{P}\in\mathcal{C}} \mathbb{E}_{\mathbb{P}}[u(f)]\,,$$

for a unique closed convex set of priors $\mathcal{C}$ (Gilboa and Schmeidler, 1989). The agent evaluates each act by its worst-case expected utility over $\mathcal{C}$. The Choquet expected-utility model (Schmeidler, 1989), integrating $u(f)$ against a non-additive "capacity," coincides with maxmin when the capacity is convex.

Two refinements trade the stark worst-case for tractable gradations. *Variational preferences* (Maccheroni et al., 2006) use $V(f) = \min_{\mathbb{P}}\{\mathbb{E}_{\mathbb{P}}[u(f)]+c(\mathbb{P})\}$, penalizing each prior by an ambiguity-index $c(\cdot)$; the *multiplier preferences* of robust control (Hansen and Sargent, 2001), with c a relative entropy from a reference model, are the leading special case and connect ambiguity to a max–min game against a malevolent "nature" that perturbs the model. The *smooth ambiguity* model (Klibanoff et al., 2005) instead writes $V(f) = \mathbb{E}_\mu[\phi(\mathbb{E}_{\mathbb{P}}[u(f)])]$ over a second-order belief μ on priors, where the concavity of ϕ *is* ambiguity aversion — separating it cleanly from risk aversion (u) much as recursive utility separates risk aversion from intertemporal substitution.

4.2 Asset-pricing consequences

Ambiguity raises required returns through a worst-case channel distinct from risk: it generates a sizeable equity premium with low effective risk aversion, produces *portfolio inertia* and *non-participation* (a region where the agent optimally does not trade), and amplifies volatility when the worst-case prior shifts with news (Epstein and Schneider, 2010). It also rationalizes *home bias* and under-diversification as responses to the ambiguity of less-familiar assets.

5 Break III — time, recursion, and customs (the intertemporal break)

5.1 Dynamic inconsistency

The additively separable discounted-utility model $\sum_t \beta^t u(c_t)$ is the intertemporal analogue of expected utility, and it breaks analogously. Strotz (1955) showed that any discount function other than exponential makes preferences *dynamically inconsistent*: the plan that is optimal today is not optimal when tomorrow arrives. The leading behavioral form, *quasi-hyperbolic* discounting with weights $\{1, \beta\delta, \beta\delta^2, \dots\}$ (Laibson, 1997), captures present bias — a disproportionate pull of immediate payoffs — and the resulting demand for commitment devices; the agent is modeled as a sequence of selves playing an intrapersonal game.

5.2 Recursive utility

A separate failure of additive EU is that it forces one parameter, γ, to govern both attitudes toward atemporal risk and toward intertemporal smoothing. Recursive utility severs them.

Definition 3 (Epstein–Zin–Weil recursion)**.** With a constant-elasticity time aggregator and a constant-relative-risk-aversion risk aggregator, lifetime utility solves

$$U_t = \Big[(1-\beta)\, c_t^{1-\rho} + \beta\, \big(\mathbb{E}_t[\, U_{t+1}^{1-\gamma}\,]\big)^{\frac{1-\rho}{1-\gamma}}\Big]^{\frac{1}{1-\rho}},$$

where γ is relative risk aversion and $\psi = 1/\rho$ is the elasticity of intertemporal substitution; EU is the knife-edge $\gamma = \rho$ (Kreps and Porteus, 1978; Epstein and Zin, 1989; Weil, 1989).

Disentangling γ from ψ lets the SDF respond to news about *future* utility, not only current consumption. This is the engine of the *long-run-risk* model (Bansal and Yaron, 2004): a small, persistent, heteroskedastic component in consumption growth, priced by a recursive-utility agent who cares about the timing of the resolution of uncertainty, generates a large equity premium and return predictability from smooth observed consumption.

5.3 Habits as the formalization of "custom"

The most direct formalization of *custom* — that satisfaction is judged against an accustomed standard, not an absolute level — is *habit formation*: utility depends on consumption relative to a slow-moving habit stock x_t, $u(c_t, x_t)$, with u increasing in c and decreasing in x. The habit may be *internal* (one's own past consumption; Constantinides, 1990) or *external* — "keeping up with the Joneses," the consumption of one's reference community (Abel, 1990). Campbell and Cochrane (1999) build a model in which the *surplus consumption ratio* $s_t = (c_t - x_t)/c_t$ drives a time-varying, countercyclical effective risk aversion γ/s_t: when consumption falls toward the customary level in a recession, risk aversion spikes, risk premia rise, and the SDF becomes volatile. Habit thus resolves the equity-premium puzzle and generates countercyclical expected returns, volatility clustering, and long-horizon predictability — from a single, economically interpretable departure of utility from absolute levels toward customary ones.

6 Break IV — frictions: liquidity and asymmetric information

The first three breaks enrich the *preferences* of a representative, fully-informed agent. The fourth is deeper: market frictions break the representativeness itself, so that no single marginal utility prices all assets.

6.1 Liquidity preference

Trading is costly, and the demand to hold readily-saleable wealth — Keynes's *liquidity preference* (Keynes, 1936) — is a first-order determinant of prices. Constantinides (1986) shows that proportional transaction costs make the relevant state a *no-trade region*: the agent rebalances only at the boundaries, and the effective SDF reflects shadow costs of trading rather than frictionless marginal utility. Amihud and Mendelson (1986) price the bid–ask spread directly: because illiquid assets are costlier to turn over, they must offer higher gross returns, a clientele effect by horizon. Liquidity is also a priced *risk*: the liquidity-adjusted CAPM of Acharya and Pedersen (2005) adds covariances between an asset's return and illiquidity and market-wide illiquidity, and Pástor and Stambaugh (2003) document a sizeable premium for sensitivity to aggregate liquidity shocks. In SDF terms, liquidity enters m both as a level (a transaction-cost wedge) and as a systematic state variable that commands its own risk price.

6.2 Asymmetric information

When agents hold *different* information, the representative-agent SDF is not merely modified — its foundations dissolve, and strategic behavior replaces price-taking.

Theorem 4 (No-trade)**.** *If risk-averse agents share a common prior and it is common knowledge that allocations are ex-ante Pareto optimal, then no purely speculative trade can occur: the* *willingness of an informed party to trade reveals its information and removes the counterparty's gain (Milgrom and Stokey, 1982). Observed trade therefore requires a departure — heterogeneous priors, noise/liquidity traders, or non-common-knowledge of rationality.*

This negative result frames the field. Akerlof (1970) shows that hidden quality (adverse selection) can unravel a market entirely — the "lemons" problem — with Spence (1973)'s signaling as the partial remedy. Grossman and Stiglitz (1980) prove an *impossibility*: if prices perfectly aggregated costly private information, no one would pay to acquire it, so informationally efficient prices cannot exist in equilibrium — some noise is necessary, and it is exactly what compensates the informed. Market microstructure makes the mechanism operational: Kyle (1985) models a strategic informed trader who optimally hides in noise-trader order flow, with price impact λ measuring market depth; Glosten and Milgrom (1985) show that a competitive market maker, facing possibly informed counterparties, sets a bid–ask spread purely as *adverse-selection* compensation, even with zero processing cost. Information asymmetry thus produces illiquidity, price impact, and a wedge between price and any agent's marginal utility — a friction the preference-based program of Sections 3–5 cannot represent.

6.3 The meta-frame: bounded rationality and limited arbitrage

Underlying all four breaks is a relaxation of perfect rationality: Simon (1955)'s *bounded rationality* (optimization is costly, so agents satisfice with heuristics); the *noise-trader* risk of De Long et al. (1990) and the *limits to arbitrage* of Shleifer and Vishny (1997), which explain why rational agents do not simply trade these features away; and the *adaptive-markets* synthesis of Lo (2004), in which preferences and degrees of rationality are an evolving population rather than fixed primitives.

7 Synthesis: a map of the frontier

The frontier is unified by the pricing kernel (1): each technology is a specific modification of m, validated by the puzzle it resolves.

Break	Preference technology	Priced consequence
Independence (Allais)	rank-dependent / prospect theory / disappointment aversion	first-order risk aversion; premium; skewness pricing
Single prior (Ellsberg)	maxmin / variational / smooth ambiguity; robust control	ambiguity premium; non-participation; home bias
Time separability	recursive utility; (quasi-)hyperbolic discounting	risk-aversion/EIS split; long-run risk; commitment demand
Absolute levels (custom)	internal/external habit formation	countercyclical risk premia; predictability; volatility
Frictionless trading	transaction costs; liquidity-risk pricing	illiquidity premium; no-trade regions
Symmetric information	adverse selection; strategic trading	bid–ask spreads; price impact; no-trade theorem

Two cross-cutting principles recur. First, several breaks deliver *first-order* (non-vanishing) risk aversion — through a kink (prospect theory), a worst case (ambiguity), or a moving reference (habit) — which is what the Hansen–Jagannathan bound (2) demands and smooth CRRA curvature cannot supply. Second, the most successful models *separate* conflated objects: risk aversion from intertemporal substitution (recursive utility), risk from ambiguity (smooth model), and the level of consumption from its customary standard (habit).

8 The open frontier: identifying preferences

1. **Identification and observational equivalence.** The deepest problem is that several distinct utilities — habit, long-run risk, rare disasters (Rietz, 1988; Barro, 2006), ambiguity — can match the same moments. Telling them apart requires over-identifying restrictions (options, higher moments, micro data on beliefs) rather than fit to the equity premium alone.

2. **Unification.** Real investors are plausibly loss-averse *and* ambiguity-averse *and* habit-forming *and* liquidity-constrained at once. Tractable preferences that combine the breaks — and equilibria with heterogeneous such agents — are largely unbuilt.

3. **Dynamic consistency.** Many non-EU and ambiguity models are dynamically inconsistent or fail recursivity; characterizing which preferences admit a well-posed sequential problem (and how a sophisticated agent should behave when they do not) remains active.

4. **Estimation from data.** Recovering preference parameters — or the SDF nonparametrically — under heavy tails and dependence is a hard statistical problem; machine-learning estimators of the pricing kernel and of belief distortions are an emerging, cross-disciplinary direction.

5. **Foundations and welfare.** If choices reveal biased preferences, which utility is the *normative* one for advice and welfare? Reference-dependent and present-biased models force an explicit split between "decision" and "experienced" utility that the classical theory never had to make.

9 Conclusion

The expected-utility core is not so much wrong as *incomplete*: it is the linear, single-prior, time-separable, frictionless, symmetric-information limit of a much larger space of preferences. The frontier of financial modeling is the disciplined exploration of departures from that limit — weakening independence for real risk attitudes, the single prior for ambiguity, time separability for recursion and habit, and the frictionless representative agent for liquidity and information. The discipline throughout is the same: every departure is a precise change to the pricing kernel of equation (1), and every change earns its place only by the asset-pricing facts it explains without doing violence to the others. The unsolved core is identification — deciding, from data, *which* departure is at work — and it is there that the next advances, statistical as much as economic, will be made.

References

Andrew B. Abel. Asset prices under habit formation and catching up with the joneses. *American Economic Review*, 80(2):38–42, 1990.

Viral V. Acharya and Lasse Heje Pedersen. Asset pricing with liquidity risk. *Journal of Financial Economics*, 77(2):375–410, 2005.

George A. Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3):488–500, 1970.

Maurice Allais. Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica*, 21(4):503–546, 1953.

Yakov Amihud and Haim Mendelson. Asset pricing and the bid-ask spread. *Journal of Financial Economics*, 17(2):223–249, 1986.

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

Nicholas Barberis and Ming Huang. Stocks as lotteries: The implications of probability weighting for security prices. *American Economic Review*, 98(5):2066–2100, 2008.

Nicholas Barberis, Ming Huang, and Tano Santos. Prospect theory and asset prices. *Quarterly Journal of Economics*, 116(1):1–53, 2001.

Robert J. Barro. Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, 121(3):823–866, 2006.

Shlomo Benartzi and Richard H. Thaler. Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics*, 110(1):73–92, 1995.

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

John H. Cochrane. *Asset Pricing*. Princeton University Press, Princeton, NJ, revised edition, 2005.

George M. Constantinides. Capital market equilibrium with transaction costs. *Journal of Political Economy*, 94(4):842–862, 1986.

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

J. Bradford De Long, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. Noise trader risk in financial markets. *Journal of Political Economy*, 98(4):703–738, 1990.

Daniel Ellsberg. Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics*, 75(4): 643–669, 1961.

Larry G. Epstein and Martin Schneider. Ambiguity and asset markets. *Annual Review of Financial Economics*, 2:315–346, 2010.

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

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.

Lawrence R. Glosten and Paul R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100, 1985.

Sanford J. Grossman and Joseph E. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408, 1980.

Faruk Gul. A theory of disappointment aversion. *Econometrica*, 59(3):667–686, 1991.

Lars Peter Hansen and Ravi Jagannathan. Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):225–262, 1991.

Lars Peter Hansen and Thomas J. Sargent. Robust control and model uncertainty. *American Economic Review*, 91(2):60–66, 2001.

Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.

John Maynard Keynes. *The General Theory of Employment, Interest and Money*. Macmillan, London, 1936.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892, 2005.

Botond Kőszegi and Matthew Rabin. A model of reference-dependent preferences. *Quarterly Journal of Economics*, 121(4):1133–1165, 2006.

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

Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.

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

Andrew W. Lo. The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management*, 30(5):15–29, 2004.

Graham Loomes and Robert Sugden. Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, 92(368):805–824, 1982.

Fabio Maccheroni, Massimo Marinacci, and Aldo Rustichini. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica*, 74(6):1447–1498, 2006.

Mark J. Machina. "expected utility" analysis without the independence axiom. *Econometrica*, 50(2):277–323, 1982.

Rajnish Mehra and Edward C. Prescott. The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161, 1985.

Paul Milgrom and Nancy Stokey. Information, trade and common knowledge. *Journal of Economic Theory*, 26(1):17–27, 1982.

Ľuboš Pástor and Robert F. Stambaugh. Liquidity risk and expected stock returns. *Journal of Political Economy*, 111(3):642–685, 2003.

John W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32(1/2):122–136, 1964.

John Quiggin. A theory of anticipated utility. *Journal of Economic Behavior & Organization*, 3 (4):323–343, 1982.

Matthew Rabin. Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, 68(5):1281–1292, 2000.

Thomas A. Rietz. The equity risk premium: A solution. *Journal of Monetary Economics*, 22(1): 117–131, 1988.

Leonard J. Savage. *The Foundations of Statistics*. John Wiley & Sons, New York, 1954.

David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57(3):571–587, 1989.

Andrei Shleifer and Robert W. Vishny. The limits of arbitrage. *Journal of Finance*, 52(1):35–55, 1997.

Herbert A. Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, 69 (1):99–118, 1955.

Michael Spence. Job market signaling. *Quarterly Journal of Economics*, 87(3):355–374, 1973.

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

Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, 1992.

John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1944.

Philippe Weil. The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics*, 24(3):401–421, 1989.

Menahem E. Yaari. The dual theory of choice under risk. *Econometrica*, 55(1):95–115, 1987.