



# Pricing Kernel Monotonicity and the Conservativeness of Risk-Neutral Forecasts

May 29, 2026

Mamiko Yamashita

Lecturer, Osaka School of International Public Policy, The University of Osaka

**【 Keywords 】** Asset Pricing, Risk Management, Pricing Kernel Puzzle, Stochastic Dominance

**【 JEL classification 】** G12; G13; G32

**【 Abstract 】** This paper studies the relationship between option-implied, risk-neutral forecasts and their real-world counterparts through the lens of stochastic dominance and pricing kernel monotonicity. We show that when the pricing kernel is weakly decreasing in asset payoffs, the real-world distribution first-order stochastically dominates the risk-neutral one, implying that the risk-neutral forecast is more conservative. The ordering is reversed when the pricing kernel is weakly increasing, implying that risk-neutral forecasts may be more optimistic rather than conservative. We further show that this monotonicity is closely linked to the dependence between asset payoffs and aggregate consumption. Our results provide a new perspective on the pricing kernel puzzle, an empirical finding that pricing kernels for major market indices are often non-monotonic. Our results, together with the pricing kernel puzzle, suggest that the commonly held belief in the conservativeness of risk-neutral forecasts is not generally warranted, even for broad market indices.

[yamashita.mamiko.osipp@osaka-u.ac.jp](mailto:yamashita.mamiko.osipp@osaka-u.ac.jp)

This is a preprint version of the paper forthcoming in *Journal of Risk*. A link to the published version will be added once it becomes available. This was supported by JSPS KAKENHI Grant Number 24K16397.

---

# Pricing Kernel Monotonicity and the Conservativeness of Risk-Neutral Forecasts

Mamiko Yamashita\*

Osaka School of International Public Policy, The University of Osaka

May 29, 2026

*This is a preprint version of the paper forthcoming in Journal of Risk. A link to the published version will be added once it becomes available.*

## Abstract

This paper studies the relationship between option-implied, risk-neutral forecasts and their real-world counterparts through the lens of stochastic dominance and pricing kernel monotonicity. We show that when the pricing kernel is weakly decreasing in asset payoffs, the real-world distribution first-order stochastically dominates the risk-neutral one, implying that the risk-neutral forecast is more conservative. The ordering is reversed when the pricing kernel is weakly increasing, implying that risk-neutral forecasts may be more optimistic rather than conservative. We further show that this monotonicity is closely linked to the dependence between asset payoffs and aggregate consumption. Our results provide a new perspective on the pricing kernel puzzle, an empirical finding that pricing kernels for major market indices are often non-monotonic. Our results, together with the pricing kernel puzzle, suggest that the commonly held belief in the conservativeness of risk-neutral forecasts is not generally warranted, even for broad market indices.

**JEL classification:** G12; G13; G32

---

\*yamashita.mamiko.osipp@osaka-u.ac.jp

# 1 Introduction

Forecasts implied by option prices have attracted considerable attention because they encode forward-looking information about future asset payoffs. In particular, option-implied volatility measures such as the VIX have become central tools in empirical finance (Poon and Granger, 2003), and market sentiment recovered from option prices is studied by Chiang and Tsai (2024).

In risk management, accurate computation of risk measures such as value-at-risk (VaR) and expected shortfall (ES) is essential. The former indicates the maximum loss at a given probability level, while the latter captures the average loss beyond that threshold and has replaced VaR in Basel capital regulations (BCBS, 2019). Beyond these standard risk measures, the Foster–Hart (FH) risk measure (Foster and Hart, 2009) captures the largest fraction of wealth that can be invested in a gamble without risking bankruptcy in a repeated-gambling framework. Although these measures are typically computed using historical data (Pérignon and Smith, 2010), there have been attempts to extract information from option prices; see Aït-Sahalia and Lo (2000); Ghysels and Wang (2014); Barone-Adesi (2016) for VaR and ES, and Leiss and Nax (2018) for FH.

A key feature of option-implied forecasts is that they are constructed under the risk-neutral measure rather than the real-world measure. This raises a fundamental question: how should risk-neutral forecasts be interpreted from the perspective of real-world risk assessment? The literature offers at least three main interpretations. First, Aït-Sahalia and Lo (2000) define economic VaR (E–VaR) as VaR under the risk-neutral measure, emphasizing its dependence on prevailing economic conditions. Second, Barone-Adesi (2016) argue that, for short horizons, risk-neutral quantities provide good approximations to their real-world counterparts. Third, Leiss and Nax (2018) show that the risk-neutral VaR of the market portfolio provides a conservative bound for the real-world VaR.

Related to the third interpretation, there is a widespread intuition that risk-neutral forecasts are “conservative,” in the sense that they assign relatively more weight to adverse outcomes.<sup>1</sup> Empirical evidence supports this view. Aït-Sahalia and Lo (2000) document that the risk-neutral density of S&P 500 returns is more negatively skewed

---

<sup>1</sup>For instance, Hatfield (2009) notes that “a risk neutral scenario set has many more really bad scenarios than a real world scenario set precisely because investors fear these scenarios.”

than the real-world density, and Bakshi et al. (2003) report similar findings for individual stocks. Moreover, Carr and Wu (2009) and Bollerslev et al. (2009) show that the variance risk premium is typically positive, implying higher variance under the risk-neutral measure. From a theoretical perspective, Leiss and Nax (2018) analyze this issue using stochastic dominance and show that, in a single-asset consumption-based model with a risk-averse representative agent and CRRA preferences, the real-world measure first-order stochastically dominates the risk-neutral measure.

Despite its widespread use, the notion that risk-neutral forecasts are inherently “conservative” lacks a clear theoretical foundation, and the economic conditions under which this interpretation is valid remain poorly understood. To address this gap, this paper provides a general theoretical characterization of when such a conservativeness interpretation is justified in terms of first-order stochastic dominance (FOSD). We link this property to the monotonicity of the projected pricing kernel (Rosenberg and Engle, 2002; Beare and Schmidt, 2016), defined as the conditional expectation of the pricing kernel given the terminal payoff of the asset or portfolio. We show that a simple monotonicity condition is sufficient for FOSD: if the projected pricing kernel is weakly decreasing in the terminal payoff, then the real-world measure first-order stochastically dominates the risk-neutral measure; if it is weakly increasing, the ordering is reversed.

Monotonicity of the projected pricing kernel, particularly for major equity indices, has been extensively studied empirically. Since the pricing kernel represents the intertemporal marginal rate of substitution, theory predicts that it should be weakly decreasing under risk aversion.<sup>2</sup> However, a large empirical literature documents violations of this prediction: estimated pricing kernels are often non-monotonic (Jackwerth, 2015; Aït-Sahalia and Lo, 2000; Rosenberg and Engle, 2002) or even U-shaped (Ziegler, 2007). This phenomenon, known as the “pricing kernel puzzle,” has been shown to be statistically significant by Beare and Schmidt (2016). See Cuesdeanu and Jackwerth (2018) for a survey.<sup>3</sup>

---

<sup>2</sup>Dybvig (1988) show that if the projected pricing kernel onto the market portfolio, denoted by  $y_t$ , is not weakly decreasing, then there exists a payoff function  $g_t$  such that  $g_t(y_{t+1})$  and  $y_{t+1}$  share the same real-world distribution conditional on time- $t$  information, yet  $g_t(y_{t+1})$  is priced lower than  $y_{t+1}$ .

<sup>3</sup>See, for example, Hens and Reichlin (2013); Siddiqi and Quiggin (2019) for explanations of the pricing kernel puzzle.

By linking pricing kernel monotonicity to FOSD, our analysis provides a new interpretation of the pricing kernel puzzle. Rather than being merely a discrepancy between theory and data, it also calls into question the common belief that risk-neutral forecasts are conservative. Even for broad indices such as the S&P 500, such an interpretation is not automatically justified. Our main contribution is to establish such a theoretical benchmark that serves as a warning against a blanket interpretation of risk-neutral forecasts as conservative.

Our framework applies beyond equity indices to assets such as commodities, exchange rates, and bonds. For example, Christoffersen et al. (2022) estimate the projected pricing kernel for crude oil and find it to be U-shaped, implying no clear evidence of either conservativeness or optimism. Beber and Brandt (2006) find a U-shaped option-implied state-price density for bonds, and Li and Zhao (2009) report similar results for LIBOR rates. Kitsul and Wright (2013) study inflation, and Bakshi et al. (2008) examine currency options. These studies consistently find non-monotonic pricing kernels, suggesting no systematic evidence of conservativeness in risk-neutral forecasts.

As a benchmark illustration, we consider a single-asset economy in which the asset payoff coincides with aggregate consumption, as in Leiss and Nax (2018), and show that the pricing kernel is weakly decreasing under risk aversion and time-separable utility. While Leiss and Nax (2018) derive their FOSD result under specific utility specifications, our analysis shows that their result does not rely on particular parametric forms such as CRRA or CARA. In a more realistic multi-asset setting, we derive a sufficient condition for monotonicity based on the dependence between asset payoffs and aggregate consumption. Although this condition is difficult to verify empirically, it provides conceptual guidance for interpreting an economic condition for a given asset.

In addition, we show how the FOSD relation implied by our framework translates into the behavior of commonly used risk measures, including Value-at-Risk (VaR), Expected Shortfall (ES), and the Foster–Hart risk measure. While these implications are straightforward, they help clarify the economic meaning of the theoretical results in the context of risk management.

Overall, this paper provides a unified theoretical framework for interpreting pricing kernel monotonicity. We do not propose a new method to recover the real-world distribution, nor do we offer a general justification for the conservativeness of risk-neutral fore-

casts. Instead, we provide a theoretical benchmark for the interpretation of risk-neutral forecasts, and for evaluating empirical findings in the pricing kernel literature.

The remainder of the paper is organized as follows. Section 2 presents the general result linking FOSD and pricing kernel monotonicity. Section 3 derives sufficient conditions within a consumption-based framework. Section 4 discusses implications for risk measurement. Section 5 concludes.

## 2 Pricing Kernel Monotonicity and Stochastic Dominance

We consider an arbitrage-free financial market with multiple traded assets, including a risk-free asset. Time is discrete and indexed by  $t = \dots, -2, -1, 0, 1, 2, \dots$ , measured in trading days. Consider an investor holding a traded asset or portfolio with price process  $\{S_t\}_{t=-\infty}^{\infty}$ . For expositional simplicity, we assume that the asset or portfolio pays no dividends, so that its payoff at time  $t + \tau$  is given by  $S_{t+\tau}$ , where  $\tau \geq 1$  denotes the forecast horizon.

Under the law of one price and the absence of arbitrage, there exists a strictly positive pricing kernel  $\{m_t^*\}_{t=-\infty}^{\infty}$  such that the pricing relation for the payoff  $S_{t+\tau}$  satisfies

$$S_t = \mathbb{E}_t^{\mathbb{P}} [m_{t:t+\tau}^* S_{t+\tau}] = \exp(-r_t \tau / 252) \mathbb{E}_t^{\mathbb{Q}} [S_{t+\tau}],$$

where  $m_{t:t+\tau}^* = m_{t+\tau}^* / m_t^*$  denotes the pricing kernel from  $t$  to  $t + \tau$ ,  $r_t$  is the annualized risk-free rate, and  $\mathbb{E}_t^{\mathbb{P}}$  and  $\mathbb{E}_t^{\mathbb{Q}}$  denote conditional expectations at time  $t$  under the real-world and risk-neutral measures, respectively. We assume that there are 252 trading days in a year. Throughout the paper, we assume that conditional on the information available at time  $t$ , the payoff  $S_{t+\tau}$  admits conditional probability density functions  $p_{t\tau}$  and  $q_{t\tau}$  under the real-world and risk-neutral measures, respectively.

It is useful to introduce the *projected pricing kernel* from  $t$  to  $t + \tau$  (Rosenberg and Engle, 2002).<sup>4</sup> The projected pricing kernel, denoted by  $m_{t:t+\tau}$ , is defined as the conditional

---

<sup>4</sup>Beare and Schmidt (2016) call  $m_{t:t+\tau}^*$  the stochastic discount factor and  $m_{t:t+\tau}$  the pricing kernel. In this paper, we follow Rosenberg and Engle (2002) and use the term “projected pricing kernel” to refer to  $m_{t:t+\tau}$ .

expectation of the intertemporal pricing kernel given the terminal payoff:

$$m_{t:t+\tau}(s) = \mathbb{E}_t^{\mathbb{P}}[m_{t:t+\tau}^* | S_{t+\tau} = s].$$

A key property of the projected pricing kernel is that it preserves the pricing implications of the original pricing kernel, in the sense that  $S_t = \mathbb{E}_t^{\mathbb{P}}[m_{t:t+\tau}(S_{t+\tau})S_{t+\tau}] = \exp(-r_t\tau/252) \mathbb{E}_t^{\mathbb{Q}}[S_{t+\tau}]$  holds. Because we do not assume a complete market, there may exist multiple pricing kernels, but the law of one price implies that the projected pricing kernel is unique. From the above equation, the real-world and risk-neutral densities are related by the equation below:

$$p_{t\tau}(s) = \frac{q_{t\tau}(s)/m_{t:t+\tau}(s)}{\int q_{t\tau}(x)/m_{t:t+\tau}(x) dx}. \quad (1)$$

Equation (1) shows that the likelihood ratio between the two measures is determined by the behavior of the projected pricing kernel. The following proposition establishes a simple sufficient condition linking the monotonicity of the pricing kernel to stochastic dominance.

**Proposition 1.** *Let  $S_t$  denote the price of an asset or portfolio at time  $t$  in an arbitrage-free market with multiple traded assets, and let  $\tau \geq 1$  denote the forecast horizon.*

(i) *Suppose that the projected pricing kernel  $m_{t:t+\tau}(S_{t+\tau}) = \mathbb{E}_t^{\mathbb{P}}(m_{t:t+\tau}^* | S_{t+\tau})$  is weakly decreasing in  $S_{t+\tau}$ . Then, the real-world measure of the terminal payoff first-order stochastically dominates the risk-neutral measure, i.e., for all  $x \in \mathbb{R}$ ,*

$$\mathbb{P}_t(S_{t+\tau} \leq x) \leq \mathbb{Q}_t(S_{t+\tau} \leq x),$$

where  $\mathbb{P}_t(\cdot)$  and  $\mathbb{Q}_t(\cdot)$  denote conditional probabilities under the real-world and risk-neutral measures, respectively, given the information available at time  $t$ .

(ii) *Suppose instead that the projected pricing kernel  $m_{t:t+\tau}(S_{t+\tau})$  is weakly increasing in  $S_{t+\tau}$ . Then, the risk-neutral measure of the terminal payoff first-order stochastically dominates the real-world measure, that is, for all  $x \in \mathbb{R}$ ,*

$$\mathbb{P}_t(S_{t+\tau} \leq x) \geq \mathbb{Q}_t(S_{t+\tau} \leq x).$$

*Proof.* From equation (1),  $\mathbb{P}_t[S_{t+\tau} \leq x] = (\int_{-\infty}^x q_{t\tau}(s)/m_{t:t+\tau}(s)ds)/(\int_{-\infty}^{\infty} q_{t\tau}(s)/m_{t:t+\tau}(s)ds)$  holds for all  $x$ . Then,

$$\begin{aligned} \mathbb{P}_t[S_{t+\tau} \leq x] - \mathbb{Q}_t[S_{t+\tau} \leq x] &= \frac{\int_{-\infty}^x q_{t\tau}(s)/m_{t:t+\tau}(s)ds}{\int_{-\infty}^{\infty} q_{t\tau}(s)/m_{t:t+\tau}(s)ds} - \int_{-\infty}^x q_{t\tau}(s)ds \\ &\propto \int_{-\infty}^x q_{t\tau}(s)/m_{t:t+\tau}(s)ds \\ &\quad - \left( \int_{-\infty}^x q_{t\tau}(s)ds \right) \left( \int_{-\infty}^{\infty} q_{t\tau}(s)/m_{t:t+\tau}(s)ds \right) \\ &= \mathbb{E}_t^{\mathbb{Q}} \left[ \frac{1}{m_{t:t+\tau}(S_{t+\tau})} \mathbb{1}_{\{S_{t+\tau} \leq x\}} \right] \\ &\quad - \mathbb{E}_t^{\mathbb{Q}} [\mathbb{1}_{\{S_{t+\tau} \leq x\}}] \mathbb{E}_t^{\mathbb{Q}} \left[ \frac{1}{m_{t:t+\tau}(S_{t+\tau})} \right] \\ &= \text{Cov}_t^{\mathbb{Q}} \left( \frac{1}{m_{t:t+\tau}(S_{t+\tau})}, \mathbb{1}_{\{S_{t+\tau} \leq x\}} \right). \end{aligned}$$

The last expression is the covariance under the risk-neutral measure between two functions of the same random variable  $S_{t+\tau}$ . Since the denominator is positive, the sign of the covariance determines the sign of  $\mathbb{P}_t[S_{t+\tau} \leq x] - \mathbb{Q}_t[S_{t+\tau} \leq x]$ . The indicator  $\mathbb{1}_{\{S_{t+\tau} \leq x\}}$  is weakly decreasing in  $S_{t+\tau}$ .

(i) If  $m_{t:t+\tau}(S_{t+\tau})$  is weakly decreasing in  $S_{t+\tau}$ , then  $m_{t:t+\tau}(S_{t+\tau})^{-1}$  is weakly increasing, so the covariance is non-positive. Hence  $\mathbb{P}_t(S_{t+\tau} \leq x) \leq \mathbb{Q}_t(S_{t+\tau} \leq x)$  for all  $x$ .

(ii) If instead  $m_{t:t+\tau}(S_{t+\tau})$  is weakly increasing in  $S_{t+\tau}$ , then  $m_{t:t+\tau}(S_{t+\tau})^{-1}$  is weakly decreasing, so the covariance is non-negative. Hence  $\mathbb{P}_t(S_{t+\tau} \leq x) \geq \mathbb{Q}_t(S_{t+\tau} \leq x)$  for all  $x$ .  $\square$

Proposition 1 establishes the sufficient condition for FOSD relation between the real-world and the risk-neutral measures of  $S_{t+\tau}$ . It is worth emphasizing that the weak monotonicity of the projected pricing kernel is a sufficient, but not necessary, condition for stochastic dominance. The dominance result is governed by the sign of the covariance between  $1/m_{t:t+\tau}(S_{t+\tau})$  and the indicator function  $\mathbf{1}_{\{S_{t+\tau} \leq x\}}$  under the risk-neutral measure.<sup>5</sup> While weak monotonicity guarantees that this covariance is non-positive for all

---

<sup>5</sup>While the dominance relation is governed by the sign of the covariance term, we focus on the monotonicity of the pricing kernel in this paper because it admits a direct economic interpretation and connects naturally to the pricing kernel literature.

$x$ , the converse does not generally hold. Stochastic dominance may therefore arise from more intricate interactions between the pricing kernel and the payoff distribution, even when the projected pricing kernel is not weakly monotonic in the payoff. Still, we focus on the monotonicity of the projected pricing kernel because it admits a clear economic interpretation and connects naturally to the pricing kernel literature, even though it is only a sufficient condition.

While this result is mathematically related to the monotone likelihood ratio (MLR) property (Milgrom, 1981), it reinterprets MLR in economically meaningful terms by mapping a statistical ordering into a condition on the pricing kernel, a central object in asset pricing. Proposition 1 provides a direct link between two strands of literature: asset pricing and risk management. On the one hand (asset pricing literature), the monotonicity of the pricing kernel has been extensively studied and is known to fail empirically (the pricing kernel puzzle); see Jackwerth (2015); Ait-Sahalia and Lo (2000); Rosenberg and Engle (2002); Ziegler (2007) for indices and Beber and Brandt (2006); Li and Zhao (2009); Christoffersen et al. (2022) for other assets. On the other hand (risk management literature), it is often implicitly assumed that risk-neutral forecasts are conservative. Proposition 1 shows that these two statements are directly related, and provides a new interpretation of the pricing kernel puzzle, as a warning that the conservativeness of risk-neutral forecasts is not a self-evident property.

We do not intend to provide an operational procedure to recover the real-world measure from the risk-neutral measure. Rather, the result should be interpreted as a theoretical characterization of when risk-neutral forecasts are more conservative or more optimistic than their real-world counterparts in the FOSD sense. In this sense, the contribution of this section is not to provide a decision rule, but to clarify the conceptual link between pricing kernel behavior and the interpretation of risk-neutral forecasts.

In the next section, we consider a consumption-based equilibrium model, which is a standard framework in the asset pricing literature, and which Leiss and Nax (2018) use to justify the conservativeness of the risk-neutral risk measures. We provide sufficient conditions for the monotonicity of the projected pricing kernel, and show that the findings of Leiss and Nax (2018) do not hinge on a particular class of utility functions.

### 3 Sufficient Conditions under a Consumption-Based Model

In the previous section, we characterized a sufficient condition for the FOSD relation between the risk-neutral and the real-world measures. In this section, we introduce a consumption-based setting to further investigate the condition. The goal of this section is not to weaken assumptions nor to provide a condition that is easier to verify empirically than the condition in the last section. Rather, we express the condition in a form that is economically interpretable in terms of the dependence between payoffs and aggregate consumption. In this framework, we characterize the main driver of the monotonicity of the projected pricing kernel. In particular, we try to gain some insight into the analysis of Leiss and Nax (2018), who consider a single-asset consumption-based setting to justify the conservativeness of risk-neutral forecasts for market indices.

Let us assume an economy with a representative agent and a time-separable utility function  $u(x)$ . Then the pricing kernel is the marginal rate of substitution of the representative agent given by  $m_{t:t+\tau}^* = \beta \frac{u'(C_{t+\tau})}{u'(C_t)}$ , where  $\beta$  is the subjective discount factor and  $C_t$  denotes aggregate consumption at time  $t$ . The following proposition highlights that the relationship between the asset payoff and aggregate consumption is crucial for determining the monotonicity of the projected pricing kernel, as well as the risk aversion of the agent.

**Proposition 2** (Sufficient conditions for the monotonicity of the projected pricing kernel under a consumption-based model). *Consider a frictionless and arbitrage-free financial market with multiple traded assets in a discrete-time setting  $t = \dots, -2, -1, 0, 1, 2, \dots$ , measured in trading days. Let  $\{S_t\}_{t=-\infty}^{\infty}$  denote the price process of a particular asset or portfolio, and let  $\mathbb{P}_t(\cdot)$  denote conditional probabilities under the real-world measure. Consider the following conditions:*

- (a) *A discrete-time consumption-based equilibrium framework holds with a representative agent who consumes  $C_t$  at time  $t$  and has a time-separable utility function  $u$ .*
- (b) *The marginal utility  $u'(c)$  is weakly decreasing in consumption  $c$ .*

(c-i) For any  $\underline{s} < \bar{s}$  and any  $x \in \mathbb{R}$ ,

$$\mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \underline{s}) \geq \mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \bar{s}).$$

(c-ii) For any  $\underline{s} < \bar{s}$  and any  $x \in \mathbb{R}$ ,

$$\mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \underline{s}) \leq \mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \bar{s}).$$

If conditions (a), (b), and (c-i) hold, then the projected pricing kernel  $m_{t:t+\tau}(S_{t+\tau})$  is weakly decreasing in  $S_{t+\tau}$ . If instead conditions (a), (b), and (c-ii) hold, then  $m_{t:t+\tau}(S_{t+\tau})$  is weakly increasing in  $S_{t+\tau}$ .

*Proof.* Condition (a) implies that the pricing kernel is given by

$$m_{t:t+\tau}^* = \beta \frac{u'(C_{t+\tau})}{u'(C_t)},$$

where  $\beta$  is the subjective discount factor and  $C_t$  denotes aggregate consumption at time  $t$ . Conditions (b) and (c-i) imply that

$$\mathbb{E}_t^{\mathbb{P}}[u'(C_{t+\tau}) \mid S_{t+\tau} = \underline{s}] \geq \mathbb{E}_t^{\mathbb{P}}[u'(C_{t+\tau}) \mid S_{t+\tau} = \bar{s}].$$

Then, the projected pricing kernel  $m_{t:t+\tau}(s)$  is weakly decreasing in  $s$ . In other words,

$$\begin{aligned} \mathbb{E}_t^{\mathbb{P}}(m_{t:t+\tau}^* \mid S_{t+\tau} = \underline{s}) &= \mathbb{E}_t^{\mathbb{P}}\left(\beta \frac{u'(C_{t+\tau})}{u'(C_t)} \mid S_{t+\tau} = \underline{s}\right) \\ &\geq \mathbb{E}_t^{\mathbb{P}}\left(\beta \frac{u'(C_{t+\tau})}{u'(C_t)} \mid S_{t+\tau} = \bar{s}\right) = \mathbb{E}_t^{\mathbb{P}}(m_{t:t+\tau}^* \mid S_{t+\tau} = \bar{s}) \end{aligned}$$

holds for any  $\underline{s} < \bar{s}$ . Analogously, if conditions (a), (b), and (c-ii) hold, then the projected pricing kernel  $m_{t:t+\tau}(S_{t+\tau})$  is weakly increasing in  $S_{t+\tau}$ . □

While conditions (c-i) and (c-ii) are difficult to verify empirically, they highlight a structural dependence between asset payoffs and aggregate consumption required for monotonicity.

Condition (c-i) states that higher realizations of the asset payoff  $S_{t+\tau}$  are associated with stochastically higher realizations of aggregate consumption  $C_{t+\tau}$ . This condition implies a positive dependence between  $C_{t+\tau}$  and  $S_{t+\tau}$ , but it is stronger than mere positive correlation. While positive correlation restricts only certain moments of the joint

distribution, Condition (c-i) imposes an ordering on the entire conditional distribution of consumption given the asset payoff. Condition (c-ii) states the opposite relation between the asset payoff and the aggregate state of the economy: higher realizations of  $S_{t+\tau}$  are associated with stochastically lower realizations of aggregate consumption  $C_{t+\tau}$ .

Broad market indices such as the S&P 500 may provide natural examples of assets for which condition (c-i) is approximately satisfied, as their payoffs tend to be tightly related to aggregate economic activity. In contrast, assets primarily used for hedging purposes, such as gold, may be illustrative for the case where condition (c-ii) is more relevant, as their payoffs can be relatively stronger in adverse states. In reality, many assets and portfolios of interest in risk management are explicitly designed for hedging, diversification, convexity, or tail protection. For such payoffs, the reverse ordering in condition (c-ii), or even the failure of both conditions, may be more relevant. More generally, when asset payoffs are not monotonically aligned with aggregate consumption, neither condition (c-i) nor condition (c-ii) may hold. This highlights that the stochastic dominance relation, and hence the conservativeness of risk-neutral forecasts, is not a general property, but depends critically on the economic role of the asset.

We impose a rather strong assumption of a frictionless market. This assumption may help explain the gap between this benchmark theory and empirical evidence in the pricing kernel puzzle. Indeed, market friction is among the leading explanations for the pricing kernel puzzle (Hens and Reichlin, 2013 and Siddiqi and Quiggin, 2019). In this sense, our frictionless benchmark is a reference point against which such empirical deviations can be understood.

When a single-asset economy is assumed, the dominance result becomes particularly transparent, as shown in the following corollary.

**Corollary 1** (Single-asset economy). *We consider a frictionless and arbitrage-free financial market with a single traded asset in a discrete-time setting, whose price process is denoted by  $\{S_t\}_{t=-\infty}^{\infty}$ . In this case, condition (c-i) holds trivially. Consequently, if conditions (a) and (b) in Proposition 2 hold, the real-world measure first-order stochastically dominates the risk-neutral measure for the terminal payoff  $S_{t+\tau}$ .*

*Proof.* In a single-asset economy, the asset payoff  $S_{t+\tau}$  coincides with aggregate consump-

tion  $C_{t+\tau}$ . When  $x < \underline{s} < \bar{s}$ ,

$$\begin{aligned}\mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \underline{s}) &= \mathbb{P}_t(S_{t+\tau} \leq x \mid S_{t+\tau} = \underline{s}) = 0 \\ \mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \bar{s}) &= \mathbb{P}_t(S_{t+\tau} \leq x \mid S_{t+\tau} = \bar{s}) = 0\end{aligned}$$

When  $\underline{s} \leq x < \bar{s}$ ,

$$\begin{aligned}\mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \underline{s}) &= \mathbb{P}_t(S_{t+\tau} \leq x \mid S_{t+\tau} = \underline{s}) = 1 \\ \mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \bar{s}) &= \mathbb{P}_t(S_{t+\tau} \leq x \mid S_{t+\tau} = \bar{s}) = 0\end{aligned}$$

When  $x \geq \bar{s}$ ,

$$\begin{aligned}\mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \underline{s}) &= \mathbb{P}_t(S_{t+\tau} \leq x \mid S_{t+\tau} = \underline{s}) = 1 \\ \mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \bar{s}) &= \mathbb{P}_t(S_{t+\tau} \leq x \mid S_{t+\tau} = \bar{s}) = 1\end{aligned}$$

In summary, in all cases,  $\mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \underline{s}) \geq \mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \bar{s})$  holds.  $\square$

Although Leiss and Nax (2018) note that FOSD holds for a specific class of utility functions shown in Bliss and Panigirtzoglou (2004), Corollary 1 shows that FOSD holds for any risk-averse, time-separable utility function. This clarifies that the dominance result of Leiss and Nax (2018) is driven primarily by the single-asset structure of the economy, rather than being driven by any specific parametric choice of the utility function. As long as the representative agent is risk averse, the utility function need not belong to commonly used parametric classes such as CRRA or exponential utility. Note that Corollary 1 does not intend to relax the assumption of Leiss and Nax (2018), but to disentangle the main driver of their result in this simple benchmark setting.

Then the next question arises: when we extend the setting in Corollary 1 to a multi-asset economy, what is the sufficient condition for FOSD? The next corollary characterizes one possible condition for the monotonicity. The key idea is that both aggregate consumption and the asset payoff are driven by a common latent state variable, and the asset payoff provides a noisy signal of the underlying economic state that determines aggregate consumption.

**Corollary 2** (Sufficient condition for (c-i) and (c-ii) in Proposition 2). *Suppose that there exist two independent processes: a latent state variable process  $\{Z_t\}_{t=-\infty}^{\infty}$  and a noise*

process  $\{\varepsilon_t\}_{t=-\infty}^{\infty}$ . Suppose that aggregate consumption  $C_t$  and the asset payoff of interest  $S_t$  satisfy

$$C_t = g(Z_t), \quad S_t = h(Z_t) + \varepsilon_t,$$

where  $g(\cdot)$  is a weakly increasing function. Assume further that  $\varepsilon_t$  is i.i.d. with a log-concave density, and that the conditional distribution of  $Z_{t+\tau}$  given the signal  $S_{t+\tau}$  admits a well-defined density. If  $h(\cdot)$  is weakly increasing, condition (c-i) in Proposition 2 holds. If  $h(\cdot)$  is weakly decreasing, condition (c-ii) in Proposition 2 holds.

*Proof.* Let  $F_{Z|S}(\cdot | s)$  denote the conditional distribution function of  $Z_{t+\tau}$  given  $S_{t+\tau} = s$  and the information at time  $t$ . By Bayes' rule and independence, its density function  $f_{Z|S}$  is proportional to

$$f_{Z|S}(z | s) \propto f_Z(z) f_\varepsilon(s - h(z)),$$

where  $f_Z$  is the conditional density function of  $Z_{t+\tau}$  given information at  $t$ . Since  $f_\varepsilon$  is log-concave, for  $\underline{s} < \bar{s}$ ,

$$\frac{f_\varepsilon(\bar{s} - v)}{f_\varepsilon(\underline{s} - v)}$$

is weakly increasing in  $v$ .

*Case 1:  $h(\cdot)$  is weakly increasing.*

Since  $h(z)$  is weakly increasing in  $z$ , it follows that  $\frac{f_{Z|S}(z|\bar{s})}{f_{Z|S}(z|\underline{s})} \propto \frac{f_\varepsilon(\bar{s}-h(z))}{f_\varepsilon(\underline{s}-h(z))}$  is weakly increasing in  $z$ . Hence,  $\{f_{Z|S}(z | s) : s \in \mathbb{R}\}$  satisfies the monotone likelihood ratio property (MLR) in  $z$ , which implies first-order stochastic dominance (FOSD):  $\mathbb{P}_t(Z_{t+\tau} \leq a | S_{t+\tau} = \bar{s}) \leq \mathbb{P}_t(Z_{t+\tau} \leq a | S_{t+\tau} = \underline{s})$  for any  $a \in \mathbb{R}$ .

*Case 2:  $h(\cdot)$  is weakly decreasing.*

In this case,  $h(z)$  is weakly decreasing in  $z$ , so the mapping  $z \mapsto v = h(z)$  is weakly decreasing. Since  $\frac{f_\varepsilon(\bar{s}-v)}{f_\varepsilon(\underline{s}-v)}$  is weakly increasing in  $v$ , it follows that  $\frac{f_\varepsilon(\bar{s}-h(z))}{f_\varepsilon(\underline{s}-h(z))}$  is weakly decreasing in  $z$ . Therefore,  $\frac{f_{Z|S}(z|\bar{s})}{f_{Z|S}(z|\underline{s})}$  is weakly decreasing in  $z$ , so the family  $\{f_{Z|S}(z | s)\}$  satisfies the reverse MLR property. This implies the reverse FOSD ordering:  $\mathbb{P}_t(Z_{t+\tau} \leq a | S_{t+\tau} = \bar{s}) \geq \mathbb{P}_t(Z_{t+\tau} \leq a | S_{t+\tau} = \underline{s})$  for any  $a \in \mathbb{R}$ .

Since  $g$  is weakly increasing, the event  $\{C_{t+\tau} \leq x\}$  is equivalent to  $\{Z_{t+\tau} \leq g^{-1}(x)\}$ , where  $g^{-1}(x) = \sup\{z : g(z) \leq x\}$  denotes the generalized inverse. Therefore,

$$\mathbb{P}_t(C_{t+\tau} \leq x | S_{t+\tau} = s) = \mathbb{P}_t(Z_{t+\tau} \leq g^{-1}(x) | S_{t+\tau} = s).$$

Applying the FOSD (or reverse FOSD) result with  $a = g^{-1}(x)$  yields

$$\mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \bar{s}) \leq \mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \underline{s}) \quad \text{if } h \text{ is increasing,}$$

$$\mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \bar{s}) \geq \mathbb{P}_t(C_{t+\tau} \leq x \mid S_{t+\tau} = \underline{s}) \quad \text{if } h \text{ is decreasing.}$$

These correspond to conditions (c-i) and (c-ii) in Proposition 2, respectively.  $\square$

Corollary 2 shows that conditions (c-i) and (c-ii) in Proposition 2 arise naturally when asset payoffs can be interpreted as noisy signals of an underlying economic state. The assumption of log-concavity of the noise density includes the normal distribution as a leading example. In such settings, asset payoffs do not perfectly reveal aggregate consumption but nevertheless contain informative content about latent macroeconomic conditions.

When  $h(\cdot)$  is weakly increasing, higher realizations of the signal  $S_{t+\tau}$  are associated with stochastically higher realizations of the latent state  $Z_{t+\tau}$ , and hence of consumption  $C_{t+\tau}$ . In this case, the signal is positively aligned with the underlying economic environment, and condition (c-i) holds. By contrast, when  $h(\cdot)$  is weakly decreasing, higher realizations of  $S_{t+\tau}$  are associated with stochastically lower realizations of  $Z_{t+\tau}$  and  $C_{t+\tau}$ . In this case, the signal is inversely related to the underlying economic environment, and condition (c-ii) holds. This situation may arise for assets whose payoffs are negatively exposed to aggregate conditions, for example, instruments that hedge against downturns or whose returns increase in adverse states.

Broad market indices, whose returns are driven by common macroeconomic factors but do not span the entire set of consumption goods or assets in the economy, naturally fit into this framework and can be viewed as noisy proxies for aggregate consumption. More generally, any asset whose payoff loads monotonically (either positively or negatively) on the latent economic state provides a signal that satisfies the stochastic monotonicity condition under the log-concavity assumption. Again, this theoretical benchmark is predicated on the assumption of a frictionless and arbitrage-free financial market. In practice, a vast empirical literature documents a pricing kernel puzzle phenomenon, and this phenomenon has been linked, among other explanations, to market incompleteness (Hens and Reichlin, 2013; Siddiqi and Quiggin, 2019).

In this section, we have characterized sufficient conditions for stochastic dominance between the risk-neutral and the real-world measures within a consumption-based equi-

librium framework. The purpose of this section is not to provide empirically verifiable conditions, but to clarify the economic mechanisms under which the sufficient condition in Proposition 1 may arise. In this sense, the role of the consumption-based model is to provide a structural interpretation of the monotonicity condition, rather than to deliver a directly testable restriction.

The next section shows how this stochastic dominance relationship can be interpreted in the context of risk management applications, with a focus on three risk measures: Value-at-Risk (VaR), Expected Shortfall (ES), and the Foster–Hart (FH) risk measure.

## 4 Implications for Risk Measures

This section illustrates how the FOSD implies conservative bounds for widely used risk measures. We focus on three commonly used risk measures: Value-at-Risk (VaR), Expected Shortfall (ES), and the Foster–Hart (FH) risk measure. Let the profit over the horizon  $[t, t + \tau]$  be defined as

$$X_{t,\tau} = S_{t+\tau} - S_t.$$

For a given confidence level  $\alpha \in (0, 1)$ , VaR and ES under the real-world measure are defined as

$$\begin{aligned} VaR_t^{\mathbb{P}}(\alpha) &= \inf\{x \in \mathbb{R} : \mathbb{P}_t(X_{t,\tau} \leq x) \geq \alpha\}, \\ ES_t^{\mathbb{P}}(\alpha) &= \frac{1}{\alpha} \int_0^\alpha VaR_t^{\mathbb{P}}(\gamma) d\gamma. \end{aligned}$$

We focus on small values of  $\alpha$ , so that VaR and ES typically take negative values. Thus, smaller values of VaR and ES correspond to higher risk. Following Leiss and Nax (2018), we define the Foster–Hart risk measure  $FH_t^{\mathbb{P}} \in [0, 1]$  as the solution  $\nu$  to

$$\mathbb{P}_t \left[ \log \left( 1 + \nu \frac{X_{t,\tau}}{S_t} \right) \right] = 0,$$

provided that  $\mathbb{E}_t^{\mathbb{P}}[X_{t,\tau}] > 0$  and  $\mathbb{E}_t^{\mathbb{Q}}[X_{t,\tau}] > 0$ , and that  $1 + (X_{t,\tau}/S_t)\nu > 0$  almost surely under both measures. If the solution exceeds one, we set  $FH_t^{\mathbb{P}} = 1$ . The risk-neutral counterparts of VaR, ES and FH, denoted by  $VaR_t^{\mathbb{Q}}$ ,  $ES_t^{\mathbb{Q}}$  and  $FH_t^{\mathbb{Q}}$ , are defined by replacing  $\mathbb{P}$  with  $\mathbb{Q}$ . In all cases, smaller values of risk measures indicate higher risk.

The following result shows that stochastic dominance implies conservative bounds for all three risk measures.

**Corollary 3.** *Suppose that the real-world measure  $\mathbb{P}$  first-order stochastically dominates the risk-neutral measure  $\mathbb{Q}$ . Then, for any  $\alpha \in (0, 1)$ , the following inequalities hold:*

$$VaR_t^{\mathbb{Q}}(\alpha) \leq VaR_t^{\mathbb{P}}(\alpha), \quad ES_t^{\mathbb{Q}}(\alpha) \leq ES_t^{\mathbb{P}}(\alpha), \quad FH_t^{\mathbb{Q}} \leq FH_t^{\mathbb{P}}.$$

*Suppose that the risk-neutral measure  $\mathbb{Q}$  first-order stochastically dominates the real-world measure  $\mathbb{P}$ . Then, for any  $\alpha \in (0, 1)$ , the following inequalities hold:*

$$VaR_t^{\mathbb{Q}}(\alpha) \geq VaR_t^{\mathbb{P}}(\alpha), \quad ES_t^{\mathbb{Q}}(\alpha) \geq ES_t^{\mathbb{P}}(\alpha), \quad FH_t^{\mathbb{Q}} \geq FH_t^{\mathbb{P}}.$$

*Proof.* The proof is given for the case where the real-world measure  $\mathbb{P}$  first-order stochastically dominates the risk-neutral measure  $\mathbb{Q}$ . The converse case is analogous.

**VaR and ES.** Fix  $\alpha \in (0, 1)$ . By Proposition 1, the real-world measure first-order stochastically dominates the risk-neutral measure, so that

$$\mathbb{P}_t(X_{t,\tau} \leq x) \leq \mathbb{Q}_t(X_{t,\tau} \leq x) \quad \text{for all } x \in \mathbb{R}.$$

In particular,

$$\mathbb{Q}_t(X_{t,\tau} \leq VaR_t^{\mathbb{P}}(\alpha)) \geq \alpha.$$

By the definition of Value-at-Risk under the risk-neutral measure,

$$VaR_t^{\mathbb{Q}}(\alpha) = \inf\{x \in \mathbb{R} : \mathbb{Q}_t(X_{t,\tau} \leq x) \geq \alpha\},$$

which implies  $VaR_t^{\mathbb{Q}}(\alpha) \leq VaR_t^{\mathbb{P}}(\alpha)$ . Since Expected Shortfall is the average of VaR over confidence levels up to  $\alpha$ , this ordering yields

$$ES_t^{\mathbb{Q}}(\alpha) \leq ES_t^{\mathbb{P}}(\alpha).$$

**FH.** Note first that  $\mathbb{E}_t^{\mathbb{P}}[f(0, X_{t,\tau})] = \mathbb{E}_t^{\mathbb{Q}}[f(0, X_{t,\tau})] = 0$ . Moreover,

$$\left. \frac{d}{d\nu} \mathbb{E}_t[f(\nu, X_{t,\tau})] \right|_{\nu=0} = \mathbb{E}_t \left[ \frac{X_{t,\tau}}{S_t} \right],$$

which is assumed to be strictly positive under both  $\mathbb{P}$  and  $\mathbb{Q}$ . Hence,  $\mathbb{E}_t^{\mathbb{P}}[f(\nu, X_{t,\tau})]$  and  $\mathbb{E}_t^{\mathbb{Q}}[f(\nu, X_{t,\tau})]$  are positive in a neighborhood of  $\nu = 0$ .

It therefore suffices to show that

$$\mathbb{E}_t^{\mathbb{Q}}[f(\nu, X_{t,\tau})] \leq \mathbb{E}_t^{\mathbb{P}}[f(\nu, X_{t,\tau})] \quad \text{for all } \nu \in [0, 1],$$

so that the root under  $\mathbb{Q}$  is weakly smaller than that under  $\mathbb{P}$ . For any fixed  $\nu \in [0, 1]$ , the function  $f(\nu, X_{t,\tau}) = \log(1 + \nu X_{t,\tau}/S_t)$  is weakly increasing in  $X_{t,\tau}$ . The inequality above then follows directly from FOSD of  $\mathbb{P}$  over  $\mathbb{Q}$ .  $\square$

This corollary shows that the risk-neutral measure yields conservative or optimistic bounds for VaR, ES, and FH risk measures, respectively, according to the direction of the monotonicity of the projected pricing kernel. Regarding the FH risk measure, Leiss and Nax (2018) discussed this property only at an informal level, whereas our proof provides a formal justification.

While the fact that FOSD implies conservative or optimistic bounds for risk measures is a general property that holds for any pair of probability measures, its relevance in the present context stems from Proposition 1. Once FOSD holds, the conservativeness or optimism of risk-neutral forecasts follows as a direct implication of the theoretical result. Importantly, this result does not provide a method to assess whether FOSD holds in practice, but rather clarifies its implications for commonly used risk measures.

## 5 Conclusion

This paper studies when option-implied risk-neutral forecasts can be interpreted as conservative or optimistic relative to their real-world counterparts. Our first contribution is to establish a general sufficient condition for first-order stochastic dominance (FOSD) based on the monotonicity of the projected pricing kernel. We show that if the projected pricing kernel is weakly decreasing, the real-world measure first-order stochastically dominates the risk-neutral measure, implying that risk-neutral forecasts can be interpreted as conservative. Conversely, if it is weakly increasing, the ordering is reversed and risk-neutral forecasts become optimistic benchmarks.

Our second contribution is to derive economically interpretable conditions that link asset payoffs to aggregate consumption. We show that the dominance result is primarily driven by the dependence structure between asset payoffs and aggregate consumption, rather than by specific parametric assumptions on utility functions such as CRRA or CARA. In a single-asset economy, the representative agent's risk aversion alone is sufficient to rationalize the conservativeness of risk-neutral forecasts. In a multi-asset setting,

this property depends on whether the asset co-moves systematically with the aggregate economy. For hedging or safe-haven assets that perform well in adverse states, the projected pricing kernel may increase, leading to a reversal of the stochastic ordering under which risk-neutral forecasts become optimistic rather than conservative.

These results provide a link between the “pricing kernel puzzle” in the asset pricing literature and the interpretation of risk-neutral forecasts. The non-monotonic pricing kernels documented empirically for market indices such as the S&P 500 suggest that the conservativeness of risk-neutral forecasts is not a general theoretical guarantee, but rather a property that may fail even for broad market indices in practice. In this sense, our framework serves as a diagnostic benchmark, highlighting that the common presumption of risk-neutral pessimism is not generally warranted.

A natural follow-up question is how such dominance conditions might be assessed in practice. Extending empirical analysis to projected pricing kernels of other assets and portfolios remains an important direction for future work. In particular, it is of interest to examine whether the (non-)monotonicity of the pricing kernel is stable over time and how it varies across macroeconomic states. The empirical construction of reliable proxies for the projected pricing kernel is also an important avenue for future work. While pricing kernels are typically estimated using asset returns, risk-free rates, and option prices, and methods such as Beare and Schmidt (2016) provide statistical tests for monotonicity, further methodological developments are needed to address estimation uncertainty.

## **Acknowledgements**

This work was supported by JSPS KAKENHI Grant Number 24K16397.

## **Declaration of Interest**

The author reports no conflict of interest. The author alone is responsible for the content and writing of this paper.

## References

- Aït-Sahalia, Y. and Lo, A. W. (2000). Nonparametric risk management and implied risk aversion. *Journal of Econometrics*, 94(1):9–51.
- Bakshi, G., Carr, P., and Wu, L. (2008). Stochastic risk premiums, stochastic skewness in currency options, and stochastic discount factors in international economies. *Journal of Financial Economics*, 87(1):132–156.
- Bakshi, G., Kapadia, N., and Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *The Review of Financial Studies*, 16(1):101–143.
- Barone-Adesi, G. (2016). VaR and CVaR implied in option prices. *Journal of Risk and Financial Management*, 9(1).
- BCBS (2019). Minimum Capital Requirements for Market Risk. Bank for International Settlements.
- Beare, B. K. and Schmidt, L. D. W. (2016). An empirical test of pricing kernel monotonicity. *Journal of Applied Econometrics*, 31(2):338–356.
- Beber, A. and Brandt, M. W. (2006). The effect of macroeconomic news on beliefs and preferences: Evidence from the options market. *Journal of Monetary Economics*, 53(8):1997–2039.
- Bliss, R. R. and Panigirtzoglou, N. (2004). Option-implied risk aversion estimates. *The Journal of Finance*, 59(1):407–446.
- Bollerslev, T., Tauchen, G., and Zhou, H. (2009). Expected stock returns and variance risk premia. *The Review of Financial Studies*, 22(11):4463–4492.
- Carr, P. and Wu, L. (2009). Variance risk premiums. *The Review of Financial Studies*, 22(3):1311–1341.
- Chiang, S. L. and Tsai, M. S. (2024). Analyzing market sentiment based on the option-implied distribution of stock returns. *Journal of Risk*, 26(5):49–93.

- Christoffersen, P., Jacobs, K., and Pan, X. N. (2022). The state price density implied by crude oil futures and option prices. *The Review of Financial Studies*, 35(2):1064–1103.
- Cuesdeanu, H. and Jackwerth, J. C. (2018). The pricing kernel puzzle: survey and outlook. *Annals of Finance*, 14:289–329.
- Dybvig, P. H. (1988). Distributional analysis of portfolio choice. *The Journal of Business*, 61(3):369–393.
- Foster, D. P. and Hart, S. (2009). An operational measure of riskiness. *Journal of Political Economy*, 117(5):785–814.
- Ghysels, E. and Wang, F. (2014). Moment-implied densities: Properties and applications. *Journal of Business & Economic Statistics*, 32(1):88–111.
- Hatfield, S. (2009). Risk-neutral valuation and its implications for product development. *Product Matters! Society of Actuaries Newsletter*, Issue 73.
- Hens, T. and Reichlin, C. (2013). Three solutions to the pricing kernel puzzle\*. *Review of Finance*, 17(3):1065–1098.
- Jackwerth, J. C. (2015). Recovering risk aversion from option prices and realized returns. *The Review of Financial Studies*, 13(2):433–451.
- Kitsul, Y. and Wright, J. H. (2013). The economics of options-implied inflation probability density functions. *Journal of Financial Economics*, 110(3):696–711.
- Leiss, M. and Nax, H. H. (2018). Option-implied objective measures of market risk. *Journal of Banking & Finance*, 88:241–249.
- Li, H. and Zhao, F. (2009). Nonparametric estimation of state-price densities implicit in interest rate cap prices. *The Review of Financial Studies*, 22(11):4335–4376.
- Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 12(2):380–391.
- Pérignon, C. and Smith, D. R. (2010). The level and quality of value-at-risk disclosure by commercial banks. *Journal of Banking & Finance*, 34(2):362–377.

- Poon, S.-H. and Granger, C. W. (2003). Forecasting volatility in financial markets: a review. *Journal Of Economic Literature*, 41(2):478–539.
- Rosenberg, J. V. and Engle, R. F. (2002). Empirical pricing kernels. *Journal of Financial Economics*, 64(3):341–372.
- Siddiqi, H. and Quiggin, J. (2019). The pricing kernel puzzle: A behavioral explanation. *Cogent Economics & Finance*, 7(1):1684609.
- Ziegler, A. (2007). Why does implied risk aversion smile? *The Review of Financial Studies*, 20(3):859–904.