Rare disasters, tail aversion, and asset pricing puzzles*

Gerrit Meyerheim

University of Munich

October 30, 2025

Abstract

This paper integrates tail aversion, implemented via a one-period entropic tilt, with rare disasters in a consumption-based asset pricing model with CRRA utility to jointly address the equity premium and risk-free rate puzzles. The model delivers closed-form expressions for the risk-free rate and asset moments, pushes out the Hansen-Jagannathan bound, implies a low risk-free rate via diffusion and disaster channels, and delivers natural upper and lower bounds of risk aversion. Calibrated to long-run return data and disciplined by disaster evidence, the model matches average returns, volatility, and a low real risk-free rate with very modest risk aversion.

JEL-classification: G12, E44, E43, E21, D81

Keywords: equity premium puzzle
risk-free rate puzzle
rare disasters
entropic tilt
multiplier (KL) preferences
robust control
consumption-based asset pricing

*I thank Alexander Matusche, Karl Schulz, Uwe Sunde, Mirko Wiederholt, and participants of various seminars for helpful comments. Support by the German Research Foundation through grants 491578970 and through CRC TRR 190 (project number 280092119) is gratefully acknowledged. Contact: Gerrit Meyerheim, University of Munich (LMU), Geschwister-Scholl-Platz 1, 80539 Munich, Germany, Email: gerrit.meyerheim@econ.lmu.de, Tel. +49 89 2180 1283.

1 Introduction

The size of the equity premium and the low real risk-free rate are among the most persistent puzzles in macro-finance. Standard consumption-based models with constant relative risk aversion (CRRA) utility and lognormal consumption growth imply Sharpe ratios and risk-free rates that are far removed from those observed in the data unless one adopts counterfactually high risk aversion (Mehra and Prescott, 1985, 2003; Cochrane, 2005). The literature has responded with a range of mechanisms that either raise the price of risk, depress the risk-free rate, or do both while preserving plausible preferences and endowment dynamics. Surveys such as Cochrane (2017) emphasize that, despite different microfoundations, many successful models share the common theme that the market's capacity to bear risk varies with macroeconomic conditions.

This paper proposes a simple synthesis that nests two empirically salient ingredients: (a) rare disasters in consumption in the spirit of Barro (2006, 2009), and (b) an entropic (risk-sensitive) tilt that overweights bad consumption states in the one-period stochastic discount factor (SDF), which shares the reduced-form implication of overweighting bad states with Benartzi and Thaler (1995). Formally, consumption growth combines a Gaussian diffusion with a Bernoulli disaster captured by parameter $j < 0$ occurring with probability p. The SDF equals the CRRA kernel with risk aversion γ scaled by an exponential change of measure in consumption growth with "tail aversion" θ. The resulting pricing kernel loads on diffusion through $\gamma + \theta$ and on disasters through a co-jump parameter $q > 1$ that captures valuation compression (e.g., leverage, default, discount-rate spikes) at disaster times.

The model delivers transparent analytical expressions for the risk-free rate, asset return moments, and the Sharpe ratio. The rare-disaster and tail aversion components push out the Hansen-Jagannathan bound and generate equity premia consistent with the data while maintaining modest within-period curvature γ. The risk-free rate falls through two channels: an interaction term $-\gamma\theta\sigma^2$ in diffusion (where σ^2 is the variance of consumption growth), and a jump ratio that implies a negative drag when $j < 0$ and $\theta > 0$. Intuitively, tail aversion acts like a state-contingent increase in marginal utility during downturns without altering the period utility function.

The framework is intentionally parsimonious. I maintain CRRA period utility and use the entropic tilt that generates tail aversion only for one-period pricing. A simple derivation shows how the SDF arises from a standard robust control problem that entails an exponential change of measure. This microfoundation facilitates comparisons with recursive-preference and robust-control literatures (Epstein and Zin, 1989, 1991; Hansen and Sargent, 2001, 2008; Maenhout, 2004), but does not rely on separating risk aversion and intertemporal elasticity of substitution. To make the strength of tail aversion transparent, I compute three quantitatively interpretable measures: the one-period Kullback–Leibler divergence, the implied total-variation bound, and the risk-neutral disaster probability. These statistics separate informational distortions from market prices and help gauge whether the extent of the tilt is economically meaningful yet informationally modest.

Two theoretical results pin down the range of plausible levels of risk aversion. First, if the risk-free rate is held fixed, there is an upper bound for risk aversion that the model can accommodate. This bound occurs exactly when tail aversion is set to zero. Second, imposing that equity values compress on impact in the event of disasters generates both a lower and an upper admissible level of risk aversion. The global upper bound is the smaller value of this new ceiling and the one implied by the risk-free rate. Within that admissible interval, the size of valuation compression the model needs is hump-shaped in risk aversion: it rises when risk aversion is very low and falls when risk aversion is very high. The intuition is straightforward. Holding the observed mean and variance of equity returns fixed, stronger compression reallocates risk from the smooth, diffusive part of returns to rare jumps. At the same time, matching the same risk-free rate forces tail aversion to decline as risk aversion increases, so the model's loading on the diffusion component first falls and then rises. Taken together, the required compression therefore reaches a unique interior maximum.

To explore the quantitative implications of the model, I calibrate preference and consumption parameters to standard values and set (p, j) using disaster evidence (Barro, 2006; Barro and Ursúa, 2008a; Nakamura, Steinsson, Barro, and Ursúa, 2013). Three parameters are then inferred from three moments: θ from the risk-free rate, co-jump q and the volatility of the diffusion part of asset return σ_r from the mean and variance of equity returns, and the drift of the diffusion part of the asset return μ_r from the asset pricing equation. The mapping from moments to parameters is analytic and exploits the mixture-lognormal structure of returns. The baseline fit reproduces average market returns and volatility as well as the low real risk-free rate found in the data by construction, implying realistic Sharpe ratios. The calibration reconciles the puzzles with modest risk aversion. The implied one-period Kullback–Leibler divergence, total-variation, and risk-neutral disaster probability show that the model is informationally hard to distinguish from a standard CRRA model while creating substantial tail aversion.

Three simple exercises help clarifying mechanics. First, shutting down either the tilt ($\theta = 0$) or disasters ($p = 0$) reveals that both channels are needed to match the risk-free rate and premia without implausible parameters. Second, sensitivity analyses show that premia reallocate between the diffusion covariance and disaster co-jumps. Raising the diffusion covariance reduces the required co-jump size, while lowering the diffusion covariance raises it. These patterns echo broader macro-finance findings on time variation in the price of risk (Cochrane, 2017) and links to state variables such as the consumption-wealth ratio (Lettau and Ludvigson, 2001). Last, I also examine how the calibration responds quantitatively to changes in γ by recomputing $\theta(\gamma)$ from the risk-free rate equation and then solving jointly for $(q(\gamma), \mu_r(\gamma), \sigma_r(\gamma))$ using the return moments and the asset pricing equation. Two robust patterns emerge. First, holding the risk-free rate fixed, $\theta(\gamma)$ is decreasing in γ in the empirically relevant range. Lower γ necessitates larger tail aversion to depress the risk-free rate, while higher γ requires smaller tail aversion. Second, to keep the equity mean and variance fixed, q adjusts to reallocate risk between diffusion and disasters. In line with the theoretical results, both higher and lower γ permit a smaller q with almost no change in σ_r. In the baseline calibration, the admissible range for risk aversion is $\gamma \in (0.46, 7.08)$. Thus, incorporating tail aversion generates a natural range for risk aversion that is quantitatively consistent with empirically observed values.

This paper makes several novel contributions. First, it offers a compact SDF that integrates rare disasters and an entropic tilt while maintaining CRRA period utility, with a clean microfoundation from a risk-sensitive recursion. Second, it is able to derive closed-form expressions for key moments (risk-free rate, Sharpe ratio, HJ bound) and an analytic inversion that maps moments to parameters (q, μ_r, σ_r). Third, it presents a quantitative calibration disciplined by disaster evidence that reconciles the equity premium and risk-free rate puzzles with potentially very low γ, even $\gamma < 1$ found in some microeconomic settings (e.g., Holt and Laury, 2002; Chetty, 2006; Cohen and Einav, 2007; Chiappori and Paiella, 2011; Barseghyan, Molinari, O'Donoghue, and Teitelbaum, 2018). Last, it provides new empirical guidance on the magnitudes of equity losses based on historical equity losses versus consumption drops. The approach builds on the rare-disaster channel (Rietz, 1988; Barro, 2006) and is complementary to habit formation (Campbell and Cochrane, 1999) and long-run risks (Bansal and Yaron, 2004). The tail aversion via entropic tilt connects to robust-control and multiplier-preference pricing (Hansen and Sargent, 2001, 2008; Maenhout, 2004) and is consistent with recursive utility foundations (Epstein and Zin, 1989, 1991). Intermediary-based models (He and Krishnamurthy, 2013; Brunnermeier and Sannikov, 2014) achieve similar macro-finance patterns via capital constraints rather than household marginal utility. In the context of the proposed model the co-jump parameter $q > 1$ can be interpreted as the valuation compression at crisis times.

The rest of this paper is organized as follows. Section 2 reviews related literature, Section 3 introduces the model and microfoundation of tail aversion, Section 4 derives analytical results and discusses economic intuition, Section 5 derives the bounds on risk aversion implied by a model with tail aversion, Section 6 presents the quantitative evaluation and sensitivity exercises, including the mapping from data to (q, μ_r, σ_r), Section 7 concludes.

2 Related literature

The equity premium puzzle highlights that representative-agent, time-additive models with plausible risk aversion cannot match the observed spread between equity and safe returns, while the companion risk-free rate puzzle emphasizes that the same models imply counterfactually high real short rates without implausible discounting or intertemporal elasticity of substitution. These facts catalyzed a large literature that proposes mechanisms to raise risk compensation, depress the risk-free rate, or do both in a disciplined way (Mehra and Prescott, 1985; Weil, 1989; Cochrane, 2005, 2017). A unifying theme is that the economy's capacity to bear risk varies over time with macroeconomic conditions, so models that embed state dependence in marginal utility (or in the representative "marginal investor") perform better in matching the data.

One prominent class of models is external habit formation (surplus consumption), which makes effective risk aversion countercyclical. When surplus consumption is low, marginal utility responds strongly to consumption, raising premia and lowering the risk-free rate. This framework matches pro-cyclical valuation ratios and return predictability (Campbell and Cochrane, 1999). A common concern is that it can imply very low and too-smooth risk-free rates and relies on sizable cyclical movements in the surplus-consumption ratio that are hard to discipline empirically (Cochrane, 2017).

A second class of models uses recursive preferences. The widely used Epstein-Zin-Weil preferences separate risk aversion from the intertemporal elasticity of substitution so that high prices of risk can co-exist with a low risk-free rate. When paired with small, persistent components in expected consumption growth and volatility, this yields the long-run risks framework, which connects premia to macro volatility and low-frequency growth movements (Epstein and Zin, 1989, 1991; Bansal and Yaron, 2004). Empirically, identification and estimation of low-frequency components are non-trivial (see Bansal, Kiku, and Yaron, 2012, for evidence and discussion).

A third class of models introduces rare disasters. A small probability of large consumption contractions increases premia and lowers the risk-free rate (Rietz, 1988; Barro, 2006). Allowing disaster intensity to vary across time produces volatile premia, return predictability, and a wide set of macro-finance regularities (Gabaix, 2012). Cross-country and long-horizon datasets provide measures for the size and frequency of disasters and the mapping from consumption drops to asset market losses (Barro and Ursúa, 2008a; Nakamura, Steinsson, Barro, and Ursúa, 2013). In practice, reconciling equity losses with consumption drops typically requires "valuation compression" at disaster times (e.g., leverage/default or higher discount rates). Martin (2008) quantifies the welfare cost of disaster risk. The co-jump parameter $q > 1$ introduced in the current paper provides a parsimonious asset-market counterpart by capturing valuation compression at disaster times. Beyond matching premia and the risk-free rate, rare-disaster models have implications for welfare costs. Barro (2009) shows that observed disaster risk can generate sizable welfare costs and helps connect the asset-pricing evidence in Barro (2006) to macro quantities. Allowing the disaster probability to vary over time improves the fit to volatility and predictability. Gabaix (2012) provides an exactly solved framework with variable disaster risk, and Wachter (2013) shows that time-varying rare-disaster risk can account for high equity volatility and a low, smooth risk-free rate.

A complementary strand emphasizes uninsurable idiosyncratic risk and incomplete markets. In Aiyagari-Bewley-Huggett environments, precautionary saving against persistent idiosyncratic income shocks raises the demand for safe assets, lowering the risk-free rate. Incorporating realistic heterogeneity and survival risk allows the stochastic discount factor of the marginal investor to be volatile even when aggregate risk is modest (Aiyagari, 1994). Building on this insight, Constantinides and Duffie (1996) show that cross-sectional heterogeneity in consumption growth and limited insurance can reconcile a high equity premium with moderate risk aversion by making the marginal investor's intertemporal marginal rate of substitution more sensitive to bad idiosyncratic outcomes. Heaton and Lucas (1996) quantify how idiosyncratic labor-income risk and limited participation reduce risk sharing and push down the risk-free rate in otherwise standard calibrations. Using panel evidence, Storesletten, Telmer, and Yaron (2004) document large and countercyclical idiosyncratic income risk over the life cycle, strengthening these mechanisms in quantitatively disciplined settings. On the cross-sectional side, the idiosyncratic volatility literature finds that higher firm-level volatility is typically not rewarded in average returns. Portfolios with high idiosyncratic volatility earn lower expected returns (Ang, Hodrick, Xing, and Zhang, 2006). Moreover, firm-level idiosyncratic volatility has increased over time (Campbell, Lettau, Malkiel, and Xu, 2001). While these facts do not directly resolve the aggregate puzzles, they reinforce the view that incomplete risk sharing and idiosyncratic tail exposure shape the supply of safe assets (risk-free rates) and the pricing of risky cash flows in bad states. This interpretation is consistent with the presented framework. As mentioned above, the co-jump parameter $q > 1$ can be read as valuation compression that is more severe for firms or intermediaries whose owners are especially exposed to idiosyncratic downside risk, while the entropic tilt magnifies the price impact of such left-tail states.

A different set of approaches uses risk-sensitive tilts and robustness. Multiplier (KL) preferences tilt probability weights toward adverse states, effectively assigning more weight to left-tail outcomes without changing within-period curvature when period utility is CRRA (Hansen and Sargent, 2001, 2008; Maenhout, 2004). In a one-period pricing kernel this appears as an exponential change of measure in consumption growth. The resulting state dependence raises the price of risk precisely when the economy is weak. The present paper adopts exactly this device via an entropic tilt and shows how it can be microfounded from a risk-sensitive recursion.

Similarly, Martin (2013) develops a cumulant-generating-function (CGF) approach to consumption-based asset pricing, highlighting that higher cumulants of consumption growth can have first-order effects on asset prices. The mixture-lognormal/disaster setting presented in this paper yields tractable CGFs, and the entropic tilt exponentially reweights the left tail of consumption growth, thereby increasing the SDF's sensitivity to left-tail cumulants while leaving period-utility curvature at γ. In this sense, the derived closed-form Sharpe ratio and HJ bound can be viewed as a CGF-based specialization with explicit jump and tilt components.

Related to these approaches, Benartzi and Thaler (1995) explain the equity premium by combining loss aversion with narrow framing (i.e. frequent evaluation). The entropic tilt presented in the current paper shares the reduced-form implication of overweighting bad states but differs in mechanism and microfoundations. Specifically, the entropic tilt is a smooth change-of-measure on consumption growth with a robust-control interpretation rather than a reference-dependent kink at a gain/loss boundary. As a result, θ raises the price of downside risk without invoking a reference point or probability weighting. For one-period (lognormal) benchmarks, an increase in evaluation frequency in the Benartzi and Thaler (1995) framework is numerically akin to a larger effective θ at short horizons, but the two approaches remain conceptually distinct. Likewise, prospect theory (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992) explains high risk premia by combining loss aversion with probability weighting that overemphasizes bad, low-probability outcomes. In contrast, the present approach generates higher disaster-state weights via the risk-neutral probabilities implied by the SDF's entropic tilt, delivering a similar pricing implication while maintaining CRRA period utility (see also Barberis, Huang, and Santos, 2001).

Finally, intermediary-based asset pricing models emphasize that the marginal investor is a constrained financial intermediary. That is, premia are high when intermediary capital is scarce and risk-bearing capacity is low (He and Krishnamurthy, 2013). Embedding a financial sector generates non-linear amplification and state-dependent risk-taking, providing a complementary resolution outside representative-household models (Brunnermeier and Sannikov, 2014).

Across these strands, the common remedy is to raise the covariance of the pricing kernel with bad states and/or lower the unconditional risk-free rate in those same states. The approach in this paper sits at the intersection of rare disasters and risk-sensitive tilting. Rare disasters contribute large negative realizations for consumption and returns. The entropic tilt increases the weight on those realizations in pricing, pushing out the Hansen-Jagannathan bound while preserving modest CRRA curvature. This combination offers a parsimonious route to reconciling the equity premium and the low risk-free rate with quantitatively disciplined parameters.

3 Model

The following section presents the model economy. After recapping the basic household problem in consumption-based asset pricing models it introduces two simple modifications to the generic model: rare disasters in endowments and returns as well as tail aversion in the preference structure.

3.1 Recap of the household problem

The generic consumption-based asset pricing problem in the spirit of Lucas Jr. (1978) and Cochrane (2005) is given by maximizing the expected discounted utility from consumption

$$\max_{\{C_s\}_{s=0}^{\infty}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t u\left(C_t\right) \right] \quad \text{with } \beta \in (0,1) \tag{1}$$

where $u(\cdot)$ is a strictly increasing, strictly concave utility function. The objective is maximized subject to the flow budget constraints

$$W_t + \phi_{t-1} X_t \geq C_t + \phi_t P_t \tag{2}$$

where W_t is an exogenous endowment, ϕ_t is the quantity of a risky asset purchased, P_t is the price of the asset and X_t is the asset's payoff. Payoffs are random variables given by the future price of the assets as well as dividend payments D_{t+1}, that is

$$X_{t+1} = P_{t+1} + D_{t+1} \tag{3}$$

The basic model can be extended to multiple assets and portfolio choice between riskless and risky assets. However, for the following analysis it suffices to assume a single risky asset and a single riskless alternative that can be bought and sold.[1] Setting up the Lagrangian of (1) subject to (2) and taking first-order conditions with respect to C_t and ϕ_t gives the core asset pricing equation

$$P_t = \mathbb{E}_t \left[M_{t+1} X_{t+1} \right] \tag{4}$$

where $M_{t+1} \equiv \beta \frac{u'(C_{t+1})}{u'(C_t)}$ is the stochastic discount factor (SDF), which ultimately links consumption growth to asset returns. That is, individuals adjust consumption today relative to tomorrow until prices and expected returns align. Define the gross return to

[1] Extending to multiple assets and portfolio choice leaves the qualitative results unaffected.

an asset as $R_{t+1} \equiv \frac{X_{t+1}}{P_t}$ with the net return defined as $R_{t+1} - 1$. This allows re-writing (4) as

$$1 = \mathbb{E}_t \left[M_{t+1} R_{t+1} \right] \tag{5}$$

As the asset is risky, a risk-free benchmark to compare asset returns to is required. If the asset were to pay out a risk-free return $R_{f,t+1}$ tomorrow then $R_{f,t+1}$ is given by re-arranging (5)

$$R_{f,t+1} = \frac{1}{\mathbb{E}_t \left[M_{t+1} \right]} \tag{6}$$

That is, a risk-free asset either pays out $X_{t+1} = R_{f,t+1}$ units of consumption tomorrow with certainty while costing one unit of consumption today or is sold at discounted price $P_t = \frac{1}{R_{f,t+1}}$ today and pays out one unit of consumption tomorrow with certainty. Using (6) and the fact that $\mathbb{E}\left[AB\right] = \mathbb{E}\left[A\right]\mathbb{E}\left[B\right] + \mathrm{Cov}\left[A, B\right]$ for two (correlated) random variables A and B allows re-arranging (5) for the risky asset to obtain the Sharpe ratio of mean excess return over standard deviation as

$$\mathrm{SR} \equiv \frac{\mathbb{E}_t \left[R_{t+1} \right] - R_{f,t+1}}{\sigma \left[R_{t+1} \right]} = - \underbrace{\frac{\mathrm{Cov}_t \left[M_{t+1}, R_{t+1} \right]}{\sigma_t \left[M_{t+1} \right] \sigma_t \left[R_{t+1} \right]}}_{\equiv \rho_{MR} \in (-1,1)} \frac{\sigma_t \left[M_{t+1} \right]}{\mathbb{E}_t \left[M_{t+1} \right]} \tag{7}$$

where ρ_{MR} is the correlation coefficient between the SDF and the asset return, $\sigma_t \left[M_{t+1} \right]$ is the standard deviation of the SDF, and $\sigma_t \left[R_{t+1} \right]$ is the standard deviation of the asset return. As the correlation coefficient between the SDF and the asset return cannot be larger than one in absolute value (7) can be simplified to

$$\left| \mathrm{SR}_{\max} \right| \leq \frac{\sigma_t \left[M_{t+1} \right]}{\mathbb{E}_t \left[M_{t+1} \right]} \tag{8}$$

which constitutes the upper envelope of the Sharpe ratio, also known as the Hansen-Jagannathan (HJ) bound (see Hansen and Jagannathan, 1991). Expressions (6) and (7) are the main moments of interest to be matched to the data. The average real annual long-run net equity return is around 7.1% while exhibiting a standard deviation of around 22% (see Jordà, Knoll, Kuvshinov, Schularick, and Taylor, 2019). Short-term bills have an average real annual long-run net return of about 1.3% in the same data. That is, the model's Sharpe ratio has to fulfill

$$\frac{7.1\% - 1.3\%}{22\%} \approx 0.26 \leq \frac{\sigma_t \left[M_{t+1} \right]}{\mathbb{E}_t \left[M_{t+1} \right]} \tag{9}$$

The "equity premium puzzle" (Mehra and Prescott, 1985) arises as the standard consumption-based asset pricing model with lognormally distributed consumption growth and CRRA utility cannot match the left-hand side of (9) by a sizable amount.

That is, for period utility and consumption growth given by

$$u\left(C_t\right)=\frac{C_t^{1-\gamma}-1}{1-\gamma} \quad \text{and} \quad \ln\frac{C_{t+1}}{C_t}\overset{i.i.d.}{\sim}\mathcal{N}\left(\mu,\sigma^2\right)$$

the moments on the right-hand side of (9) are given by

$$\mathbb{E}_t\left[M_{t+1}\right]=\beta e^{-\gamma\mu+\frac{1}{2}\gamma^2\sigma^2} \quad \text{and} \quad \sigma_t\left[M_{t+1}\right]=\beta e^{-\gamma\mu+\frac{1}{2}\gamma^2\sigma^2}\sqrt{e^{\gamma^2\sigma^2}-1}$$

where μ is mean consumption growth and σ^2 is the variance of consumption growth. This yields the log risk-free rate and HJ bound as

$$\begin{aligned}\ln R_{f,t+1}=&\delta+\gamma\mu-\frac{1}{2}\gamma^2\sigma^2\\ \left|\mathrm{SR}_{\max}\right|\approx&\gamma\sigma\end{aligned}$$

where $\delta=-\ln\beta$. Long-run growth in (non-durable) consumption per capita has a standard deviation of around 2% per year. This implies

$$0.26\leq\gamma\times 0.02 \quad\iff\quad \gamma\geq 13$$

which is an unreasonably high degree of risk aversion. Additionally, this assumes perfect correlation between the SDF and asset returns. The general case given by

$$\left|\frac{\mathbb{E}_t\left[R_{t+1}\right]-R_{f,t+1}}{\sigma_t\left[R_{t+1}\right]}\right|=\rho_{MR}\gamma\sigma$$

requires $\gamma\approx 65$ for $\rho_{MR}\approx 0.2$ (as is typically the correlation between consumption growth and asset returns). Furthermore, even if $\gamma=65$ were acceptable, the log risk-free rate is then given by

$$\ln R_{f,t+1}\approx\delta+0.78$$

for $\mu=0.025$ and $\sigma=0.02$. This in turn implies a negative rate of time preference if the model is to match the data on observed risk-free rates and constitutes the "risk-free rate puzzle" (Weil, 1989). The following parts of this section describe the structure imposed on the SDF M_{t+1} and asset return R_{t+1}. The approach combines rare disasters with an entropic tilt. Together they expand the HJ bound while maintaining modest curvature.

3.2 Endowments and returns

Let consumption growth be characterized by

$$g_{t+1} \equiv \ln \frac{C_{t+1}}{C_t} = z_{t+1} + J_{t+1} \tag{10}$$

where $z_{t+1} \overset{i.i.d.}{\sim} \mathcal{N}\left(\mu, \sigma^2\right)$ and $J_{t+1} \in \{0; j\}$ is a Bernoulli disaster jump with $j < 0$ that occurs with probability $p \in [0, 1]$, that is

$$\Pr\left(J_{t+1} = j\right) = p \tag{11}$$

which follows Barro (2006). For simplicity, the risky asset has diffusive log return $r_{t+1} \overset{i.i.d.}{\sim} \mathcal{N}\left(\mu_r, \sigma_r^2\right)$ with

$$\mathrm{Corr}(r_{t+1}, z_{t+1}) \equiv \rho$$

Assume that jumps and diffusion are uncorrelated, i.e.

$$J_{t+1} \perp\!\!\!\perp (z_{t+1}, r_{t+1})$$

However, if the consumption disaster occurs, the asset return co-jumps by $qj < 0$ with $q > 1$.[2] As a result, the (gross) asset return is given by

$$R_{t+1} = \begin{cases} e^{r_{t+1}} & \text{if } J = 0, \\ e^{r_{t+1} + qj} & \text{if } J = j \end{cases} \tag{12}$$

There are three noteworthy aspects about this specification. First, the specification of $q > 1$ implies that the price-dividend ratio moves at the consumption disaster. If the representative asset pays an aggregate dividend D_t equal to a claim on consumption (i.e. $D_t = C_t$), then in the case of a constant price-dividend ratio $P_t = \kappa D_t$ the gross return is given by

$$R_{t+1} = \frac{P_{t+1} + D_{t+1}}{P_t} = \frac{\kappa + 1}{\kappa} \frac{D_{t+1}}{D_t}$$

Hence $\ln R_{t+1} = \ln\left(\frac{\kappa+1}{\kappa}\right) + \ln \frac{D_{t+1}}{D_t}$. A disaster multiplies D_{t+1} by e^j, so $\Delta \ln R_{t+1} = j$. That is, the consumption (and hence dividend) disaster translates one-to-one into the return and thus $q \equiv 1$. If the price-dividend ratio moves at the consumption disaster (e.g., discount rates rise, risk premia spike etc.), then

$$\Delta \ln R_{t+1} = j + \Delta \ln \kappa_{t+1} \quad \text{and} \quad q = 1 + \frac{\Delta \ln \kappa_{t+1}}{j}$$

Since $\Delta \ln \kappa_{t+1} < 0$ in downturns as valuation compresses, this implies $q > 1$.[3] Sub-

[2] For further empirical evidence on this channel see Barro and Ursúa (2008a); Nakamura, Steinsson, Barro, and Ursúa (2013) among others.

[3] Appendix A contains a simple model for valuation compression using standard asset pricing arguments.

sequently, the asset price falls by more than the one-period dividend shortfall. More specifically, if a disaster cuts consumption by a fraction $x \in [0,1]$, i.e. $j = \ln(1-x)$, and the equity loss on impact is given by the fraction $L \in [0,1]$, then

$$qj = \ln(1-L) \quad \Longrightarrow \quad q = \frac{\ln(1-L)}{\ln(1-x)} \tag{13}$$

Second, even though this specification contains disastrous jumps it does not imply a fat-tailed distribution of the return. To see this note that log returns are given by $\ln R_{t+1} = r_{t+1} + q\,J_{t+1}$. The moment generating function of $\ln R_{t+1}$ is finite for all $t \in \mathbb{R}$

$$M_R(t) \equiv \mathbb{E}\left[e^{t\ln R_{t+1}}\right] = e^{t\mu_r + \frac{1}{2}t^2\sigma^2}\left(1 - p + p\,e^{tqj}\right) < \infty$$

implying that its tail probabilities decay at least exponentially for some $t > 0$. Therefore, although the model increases the chances of an "outlier" relative to a single normal distribution, the far-tail decay rate is still governed by the Gaussian diffusion. Hence, the results do not rely on a fat-tailed distribution of returns.

Third, the model adopts a one-period, on-impact disaster realization for simplicity and transparency. This choice preserves a closed-form mixture-lognormal solution for the return moments as well as a compact risk-free rate without adding a persistence state.

3.3 Preferences

Let period utility be given by standard CRRA preferences with coefficient of relative risk aversion $\gamma > 0$. Further, let the subjective discount rate be $\delta > 0$. In isolation this implies an inter-temporal marginal rate of substitution given by

$$\mathrm{IMRS}_{t,t+1} \equiv \beta\frac{u'(C_{t+1})}{u'(C_t)} = e^{-\delta}\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma} = e^{-\delta-\gamma g_{t+1}}$$

To characterize the importance of avoiding large losses in utility, let period utility be adjusted by an entropic tilt that replaces the standard conditional expectation. Let S_{t+1} be a shock (or vector of shocks) whose left-tail realizations are overweighted by the household for a given choice of future consumption C_{t+1}. The entropic tilt is then given by

$$\mathbb{E}_t^{(\theta)}\left[C_{t+1}\right] \equiv \frac{\mathbb{E}_t\left[C_{t+1}\,e^{-\theta\,S_{t+1}}\right]}{\mathbb{E}_t\left[e^{-\theta\,S_{t+1}}\right]}$$

The interpretation $q > 1$ is also consistent with the evidence that stock-price declines during disasters typically exceed the contemporaneous drop in consumption/dividends (e.g. Barro and Ursúa, 2008b) and with frameworks where disaster news raises required returns (Gabaix, 2012; Wachter, 2013). It is also compatible with the levered/unlevered distinction stressed in Barro (2006, 2009).

where $\theta > 0$ governs the degree of tail aversion. This is equivalent to a normalized exponential change of measure associated with KL ("multiplier") preferences (see Kullback and Leibler, 1951; Hansen and Sargent, 2001, 2008; Maenhout, 2004). If one restricts S_{t+1} to consumption growth, the one-period pricing operator is defined by

$$\mathbb{E}_t^{(\theta)}\left[C_{t+1}\right] \equiv \frac{\mathbb{E}_t\left[C_{t+1}\, e^{-\theta\, g_{t+1}}\right]}{\mathbb{E}_t\left[e^{-\theta\, g_{t+1}}\right]} \tag{14}$$

which implies an inter-temporal marginal rate of substitution given by

$$\text{IMRS}_{t,t+1} \equiv \underbrace{\beta \frac{u'\left(C_{t+1}\right)}{u'\left(C_t\right)}}_{\text{utility MRS}} \underbrace{\frac{e^{-\theta\, g_{t+1}}}{\mathbb{E}_t\left[e^{-\theta\, g_{t+1}}\right]}}_{\text{tail adjustment}}$$

There are two equivalent routes to microfound the tilt. First, one can axiomatize a one-period entropic change of measure that depends only on consumption growth. This yields an exact SDF immediately. Second, one can adopt a risk-sensitive recursion with multiplier (KL) preferences under full information, which yields a general SDF. When the continuation value is (locally) affine in consumption growth, the two are observationally equivalent at one period. The following Proposition describes how an exponential change of measure yields an exact tail-adjusted SDF in closed-form.

Proposition 1. *For any $\mathcal{F}_{t+1}$-measurable $X_{t+1} : \Omega \to \mathbb{R}$, let the one-period pricing operator be defined by*

$$\mathbb{E}_t^{(\theta)}\left[X_{t+1}\right] \equiv \mathbb{E}_t^{\mathbb{Q}_\theta}\left[X_{t+1}\right] = \frac{\mathbb{E}_t\left[X_{t+1}\, e^{-\theta g_{t+1}}\right]}{\mathbb{E}_t\left[e^{-\theta g_{t+1}}\right]}$$

with Radon-Nikodym derivative

$$\Lambda_{t+1} \equiv \left.\frac{d\mathbb{Q}_\theta}{d\mathbb{P}}\right|_{\mathcal{F}_{t+1}} = \frac{e^{-\theta g_{t+1}}}{\mathbb{E}_t\left[e^{-\theta g_{t+1}}\right]}$$

where $\mathbb{Q}_\theta$ is the entropic-tilt measure and $\mathbb{P}$ is the physical (i.e. data-generating) measure. Then the recursion

$$U_t = u\left(C_t\right) + \beta\, \mathbb{E}_t\left[\Lambda_{t+1}\, U_{t+1}\right]$$

yields the exact SDF

$$M_{t+1} = \beta \frac{u'\left(C_{t+1}\right)}{u'\left(C_t\right)} \frac{e^{-\theta\, g_{t+1}}}{\mathbb{E}_t\left[e^{-\theta\, g_{t+1}}\right]} = \frac{e^{-\delta-(\gamma+\theta) g_{t+1}}}{\mathbb{E}_t\left[e^{-\theta\, g_{t+1}}\right]}$$

given CRRA period utility and provided that $\mathbb{E}_t\left[e^{-\theta g_{t+1}}\right] \in (0, \infty)$ almost surely.

Proof. Follows immediately from the Radon-Nikodym derivative Λ_{t+1} on $(\Omega, \mathcal{F}_{t+1})$

and the FOC for asset demand. By construction $\mathbb{E}_t[\Lambda_{t+1}] = 1$ and $\Lambda_{t+1} \geq 0$, so $\mathbb{Q}_\theta$ is a well-defined probability kernel on $(\Omega, \mathcal{F}_{t+1})$.

$\square$

For the robust-control (risk-sensitive) recursion below I impose no measurability restriction: the worst-case density m is allowed to be $\mathcal{F}_{t+1}$-measurable (full information). The following definition lays out the robust-control operator.

Definition 1. *For any $\mathcal{F}_{t+1}$-measurable X_{t+1}, define*

$$\mathcal{T}_\theta(X_{t+1}) \equiv \inf_{m \in \mathcal{M}} \left\{ \mathbb{E}_t[m\, X_{t+1}] + \frac{1}{\theta}\mathbb{E}_t[m \ln m] \right\} \quad s.t.\ \mathbb{E}_t[m] = 1$$

where $\mathcal{M} := \{ m \geq 0 : m \text{ is } \mathcal{F}_{t+1}\text{-measurable} \}$.

The solution to this operator is summarized in the following Lemma.

Lemma 1. *Let $X_{t+1} \in L^1(\mathcal{F}_{t+1})$ and $\mathbb{E}_t\left[e^{-\theta X_{t+1}}\right] \in (0, \infty)$ almost surely, then*

$$\mathcal{T}_\theta(X_{t+1}) = -\frac{1}{\theta} \ln \mathbb{E}_t\left[e^{-\theta\, X_{t+1}}\right] \tag{15}$$

while the unique minimizer is given by

$$m^\star = \frac{e^{-\theta\, X_{t+1}}}{\mathbb{E}_t\left[e^{-\theta\, X_{t+1}}\right]} \tag{16}$$

Proof. Consider the Lagrangian given by

$$\mathcal{L}(m, \lambda) = \mathbb{E}_t\left[m X_{t+1} + \frac{1}{\theta} m \ln m - \lambda(m-1)\right]$$

where λ is the Lagrange multiplier. The integrand $m \mapsto mX + \frac{1}{\theta} m \ln m$ is strictly convex on $[0, \infty)$. As a result, the conditional first-order condition is necessary and sufficient. Differentiating pointwise and setting to zero yields

$$X_{t+1} + \frac{1}{\theta}[1 + \ln m^\star] - \lambda = 0$$

Solving for $m^\star$ gives

$$m^\star = e^{\theta\lambda - 1 - \theta\, X_{t+1}} = \frac{e^{-\theta\, X_{t+1}}}{Z}$$

with $Z = e^{1-\theta\lambda}$. Satisfying the constraint $\mathbb{E}_t[m^\star] = 1$ requires $Z = \mathbb{E}_t\left[e^{-\theta\, X_{t+1}}\right]$ and thus yields

$$m^\star = \frac{e^{-\theta\, X_{t+1}}}{\mathbb{E}_t\left[e^{-\theta\, X_{t+1}}\right]}$$

Substituting back into the objective gives the optimal value as

$$\mathcal{T}_\theta\left(X_{t+1}\right)=\mathbb{E}_t\left[m^\star X_{t+1}+\frac{1}{\theta}\,m^\star\,\ln\frac{e^{-\theta\,X_{t+1}}}{\mathbb{E}_t\left[e^{-\theta\,X_{t+1}}\right]}\right]=-\frac{1}{\theta}\ln\mathbb{E}_t\left[e^{-\theta\,X_{t+1}}\right]$$

Strict convexity of $m\,\ln m$ ensures uniqueness of the result.

□

The next Proposition summarizes how the robust control problem embedded in a risk-sensitive recursion generates a tilted SDF in general.

Proposition 2. *Consider the risk-sensitive recursion* $U_t=u\left(C_t\right)+\beta\,\mathcal{T}_\theta\left(U_{t+1}\right)$ *with* $\mathbb{E}_t\left[e^{-\theta U_{t+1}}\right]\in(0,\infty)$ *almost surely. Then the one-period SDF is given by*

$$M_{t+1}=\beta\frac{u'\left(C_{t+1}\right)}{u'\left(C_t\right)}\,\frac{e^{-\theta\,U_{t+1}}}{\mathbb{E}_t\left[e^{-\theta\,U_{t+1}}\right]} \tag{17}$$

Proof. Let P_t be the price of a payoff X_{t+1} in one period. Embed the payoff into the period-t budget and consider a marginal variation. The envelope theorem for the recursion

$$U_t=u\left(C_t\right)+\beta\inf_{m\geq 0,\mathbb{E}_t[m]=1}\left\{\mathbb{E}_t\left[mU_{t+1}\right]+\frac{1}{\theta}\,\mathbb{E}_t\left[m\ln m\right]\right\}$$

implies that the marginal value of one extra unit of payoff at $t+1$ is given by

$$\beta\,\mathbb{E}_t\left[m^\star\,u'\left(C_{t+1}\right)\,\frac{\partial C_{t+1}}{\partial X_{t+1}}\right]$$

evaluated at the optimizer $m^\star$ from Lemma 1. A one-to-one mapping between payoff and consumption implies $\frac{\partial C_{t+1}}{\partial X_{t+1}}=1$, while the marginal cost today is given by $u'\left(C_t\right)$. Hence the SDF is given by

$$M_{t+1}=\beta\frac{u'\left(C_{t+1}\right)}{u'\left(C_t\right)}\,\frac{e^{-\theta\,U_{t+1}}}{\mathbb{E}_t\left[e^{-\theta\,U_{t+1}}\right]}$$

which is equivalent to (17).

□

If one keeps the risk-sensitive recursion and projects $U_{t+1}=a_t+b_t\,g_{t+1}$, then the effective tilt is given by $\theta^\star=\theta\,b_t$ The following Lemma lays out simple sufficient conditions for this to hold.

Lemma 2. *Let the future state (or state vector) be given by*

$$S_{t+1}=\varphi S_t+\psi g_{t+1}+\epsilon_{t+1}$$

with $\mathbb{E}_t[\epsilon_{t+1}] = 0$*. If* $U_{t+1} = \chi + \tau S_{t+1}$*, then* $U_{t+1} = a_t + b_t\, g_{t+1} + \varepsilon_{t+1}$ *is affine in* g_{t+1} *with* $\mathbb{E}_t[\varepsilon_{t+1}] = 0$*.*

Proof. Follows immediately from substitution and linearity of expectations, that is $U_{t+1} = \chi + \tau\,\varphi S_t + \tau\,\psi g_{t+1} + \tau\,\epsilon_{t+1}$, with $a_t = \chi + \tau\,\varphi S_t$, $b_t = \tau\,\psi$, $\varepsilon_{t+1} = \tau\,\epsilon_{t+1}$, and $\mathbb{E}_t[\varepsilon_{t+1}] = 0$.

□

The last Proposition establishes the existence of an exact one-period SDF in this general case.

Proposition 3. *Suppose* $U_{t+1} = a_t + b_t\, g_{t+1}$*. With CRRA period utility* (17) *reduces to*

$$M_{t+1} = \frac{e^{-\delta - (\gamma + \theta\, b_t) g_{t+1}}}{\mathbb{E}_t\left[e^{-\theta\, b_t\, g_{t+1}}\right]}$$

with effective tilt $\theta^\star = \theta\, b_t$*.*

Proof. Follows from inserting $U_{t+1} = a_t + b_t\, g_{t+1}$ in (17), applying $\beta \frac{u'(C_{t+1})}{u'(C_t)} = e^{-\delta - \gamma g_{t+1}}$, and collecting terms.

□

A few aspects are noteworthy about this result. First, write the exact projection as $U_{t+1} = a_t + b_t\, g_{t+1} + \varepsilon_{t+1}$ with $\mathbb{E}_t[\varepsilon_{t+1}] = 0$. If ε_{t+1} is conditionally sub-Gaussian given $\mathcal{F}_t$ with variance proxy σ_ε^2, e.g., $\varepsilon_{t+1} | \mathcal{F}_t \sim \mathcal{N}\left(0, \sigma_\varepsilon^2\right)$, then the risk-free rate from the affine projection (without the error term) has the following projection error

$$\left|\Delta \ln R_f\right| \le \frac{1}{2}\theta^2 \mathrm{Var}_t[\varepsilon_{t+1}]$$

while the distance to the HJ bound shifts by at most $O\left(\theta^2\, \mathrm{Var}_t[\varepsilon_{t+1}]\right)$. Hence accuracy is governed by the residual variance left after projecting U_{t+1} on g_{t+1}. If the projection captures a fraction R^2 of the variation in U_{t+1} attributable to g_{t+1}, then $\mathrm{Var}_t[\varepsilon_{t+1}] = b_t^2 \frac{1-R^2}{R^2} \mathrm{Var}_t[g_{t+1}]$. It is reasonable to assume that R^2 is at least moderate as consumption growth shocks are small at annual frequency and the continuation value is a smooth functional of the states. Therefore $R^2 > 0.5$ is appropriate which implies, combined with $\mathrm{Var}_t[g_{t+1}] = (0.02)^2$, that the error is quantitatively negligible even for medium $\theta^\star$. Additionally, if ε_{t+1} is conditionally independent of g_{t+1} then

$$\frac{M_{t+1}^{\text{exact}}}{M_{t+1}^{\text{affine}}} = \frac{e^{-\theta\,\varepsilon_{t+1}}}{\mathbb{E}_t\left[e^{-\theta\,\varepsilon_{t+1}}\right]} \quad \Longrightarrow \quad \mathbb{E}_t\left[M_{t+1}^{\text{exact}}\right] = \mathbb{E}_t\left[M_{t+1}^{\text{affine}}\right]$$

and there is no projection error in the risk-free rate. Altogether, retaining the risk-sensitive recursion and projecting onto g_{t+1} yields a stochastic discount factor with tail aversion $\theta^\star = \theta b_t$ and a tightly bounded approximation error. In contrast, Proposition

1 justifies the tail-adjusted one-period SDF exactly under the $\sigma[g_{t+1}]$ entropic pricing restriction.

Second, analogous possible microfoundations include ambiguity aversion with worst-case re-weighting as well as ruin or survival constraints approximated by exponential tilting of bad tail probabilities. All approaches yield "truly" stochastic subjective discounting that is negatively correlated with consumption growth in reduced form.

Third, without uncertainty in the marginal utility generated from consumption growth the tail adjustment collapses to one in expectations. As a result, the tail-adjusted SDF is given by

$$M_{t+1} = \frac{e^{-\delta-(\gamma+\theta)g_{t+1}}}{\mathbb{E}_t\left[e^{-\theta\, g_{t+1}}\right]} \tag{18}$$

Note that the one-period SDF applies an entropic tilt $\theta > 0$ to bad consumption states (i.e. a larger θ implies more weight on low C_{t+1}) while the aggregator remains CRRA. That is, only the one-period SDF is tilted. Therefore, the local log-SDF loading on consumption growth is $-(\gamma+\theta)$, but the relative risk aversion remains γ. Put differently, the tail adjustment raises the price of consumption risk but does not alter utility curvature.

4 Analytic results

The following section presents the analytic results from the model and their economic intuition. To lighten notation, I drop time subscripts as the price is always defined for period t, the payoff is always defined for period $t+1$, and expectations are conditional on time-t information.

4.1 Risk-free rate, return moments, and equity premia

Since consumption growth is given by a lognormal component and a Bernoulli jump its expectation is given by

$$\mathbb{E}\left[e^{-\alpha g}\right] = e^{-\alpha\mu+\frac{1}{2}\alpha^2\sigma^2}\left(1-p+p\,e^{-\alpha j}\right) \tag{19}$$

Under the specified mixture-normal consumption growth, the denominator of the SDF for all $\theta > 0$ is finite by (19). Combining (6) with (18) and using (19) gives the log risk-free rate as

$$\ln R_f = \delta + \gamma\mu - \frac{1}{2}\gamma^2\sigma^2 - \gamma\,\theta\,\sigma^2 + \ln\left[\frac{1-p+p\,e^{-\theta j}}{1-p+p\,e^{-(\gamma+\theta)j}}\right] \tag{20}$$

Combining (5), (12), and (18) while using (19) implies the following for the mean log (diffusive) return implied by the asset pricing equation

$$\begin{aligned}\mu_r + \frac{1}{2}\sigma_r^2 =& \delta + \gamma\mu - \frac{1}{2}\gamma^2\sigma^2 - \gamma\,\theta\,\sigma^2 + (\gamma+\theta)\,\rho\,\sigma\,\sigma_r \\ &+ \ln\left[\frac{1-p+p\,e^{-\theta\,j}}{1-p+p\,e^{-(\gamma+\theta-q)j}}\right]\end{aligned} \tag{21}$$

Return moments are given by

$$\mathbb{E}\left[R\right] = e^{\mu_r+\frac{1}{2}\sigma_r^2}\,\left(1-p+p\,e^{qj}\right) \tag{22}$$

$$\mathbb{E}\left[R^2\right] = e^{2\mu_r+2\sigma_r^2}\,\left(1-p+p\,e^{2qj}\right)$$

$$\mathrm{Var}\left[R\right] = \mathbb{E}\left[R^2\right] - \left(\mathbb{E}\left[R\right]\right)^2 \tag{23}$$

Combining (7), (21), (22), and (23) yields the Sharpe ratio as

$$\mathrm{SR} = \frac{e^{\mu_r+\frac{1}{2}\sigma_r^2}\,\left(1-p+p\,e^{qj}\right) - e^{\delta+\gamma\mu-\frac{1}{2}\gamma^2\sigma^2-\gamma\,\theta\,\sigma^2}\,\frac{1-p+p\,e^{-\theta\,j}}{1-p+p\,e^{-(\gamma+\theta)j}}}{\sqrt{e^{2\mu_r+2\sigma_r^2}\,\left(1-p+p\,e^{2qj}\right) - \left[e^{\mu_r+\frac{1}{2}\sigma_r^2}\,\left(1-p+p\,e^{qj}\right)\right]^2}} \tag{24}$$

For small risks this can be approximated by

$$\mathrm{SR} \approx \frac{(\gamma+\theta)\,\sigma\,\sigma_r\,\rho - p\left(e^{qj}-1\right)\left(e^{-(\gamma+\theta)j}-1\right)}{\sqrt{\sigma_r^2 + p\,(1-p)\,\left(e^{qj}-1\right)^2}}$$

Finally, the HJ bound is given by using (18) and applying (19)

$$\mathrm{SR}_{\max} = \sqrt{\frac{\mathbb{E}\left[M^2\right]}{\mathbb{E}\left[M\right]^2} - 1} = \sqrt{e^{(\gamma+\theta)^2\sigma^2}\,\frac{1-p+p\,e^{-2(\gamma+\theta)j}}{\left(1-p+p\,e^{-(\gamma+\theta)j}\right)^2} - 1} \tag{25}$$

4.2 Measures of tail aversion

In order to quantify the impact of tail aversion I compute two measures of its strength: one from the preference-side, one from the market-side. By (18), the entropic-tilt measure and the one-period pricing operator are characterized by

$$\frac{d\mathbb{Q}_\theta}{d\mathbb{P}} = \frac{e^{-\theta g}}{\mathbb{E}\left[e^{-\theta g}\right]} \quad \text{and} \quad \mathbb{E}^{\mathbb{Q}_\theta}\left[X\right] = \frac{\mathbb{E}\left[X\,e^{-\theta g}\right]}{\mathbb{E}\left[e^{-\theta g}\right]}$$

Then the one-period Kullback-Leibler divergence of $\mathbb{Q}_\theta$ from $\mathbb{P}$ is given by

$$D_{\mathrm{KL}}\left(\mathbb{Q}_\theta \| \mathbb{P}\right) = \mathbb{E}^{\mathbb{Q}_\theta}\left[\ln\frac{d\mathbb{Q}_\theta}{d\mathbb{P}}\right] = \mathbb{E}^{\mathbb{Q}_\theta}\left[-\theta\,g\right] - \ln\mathbb{E}\left[e^{-\theta g}\right] \tag{26}$$

In order to express (26) in closed-form note the following equivalence

$$\mathbb{E}^{\mathbb{Q}_\theta}[g] = \frac{\mathbb{E}\left[g\, e^{-\theta g}\right]}{\mathbb{E}\left[e^{-\theta g}\right]} = -\frac{d \ln \mathbb{E}\left[e^{-\theta g}\right]}{d\theta} \tag{27}$$

Plugging (27) back into (26) while using (19) gives the one-period Kullback-Leibler divergence as

$$D_{\mathrm{KL}}\left(\mathbb{Q}_\theta \| \mathbb{P}\right) = \frac{1}{2}\theta^2 \sigma^2 - \ln\left(1 - p + p\, e^{-\theta j}\right) - \frac{p\,\theta\, j\, e^{-\theta j}}{1 - p + p\, e^{-\theta j}} \tag{28}$$

where $D_{\mathrm{KL}}\left(\mathbb{Q}_\theta \| \mathbb{P}\right)$ measures (in nats) how far the entropically tilted preference measure $\mathbb{Q}_\theta$ is from the physical measure $\mathbb{P}$. Numerically, small per-period values mean the preference distortion is informationally modest. As a rule of thumb, under independence across years it takes $\frac{1}{D_{\mathrm{KL}}(\mathbb{Q}_\theta \| \mathbb{P})}$ years in a yearly setting to distinguish between $\theta > 0$ and $\theta = 0$ from an information perspective. Pinsker's bound is given by

$$|\mathbb{Q}_\theta - \mathbb{P}|_{\mathrm{TV}} \leq \sqrt{\frac{1}{2} D_{\mathrm{KL}}} \in [0, 1] \tag{29}$$

This statistic measures the total-variation distance, which the change in probability of an event (in percentage points). Both statistics constitute a preference-side gauge of how strong the tail aversion is.

An orthogonal measure is given by risk-neutral probabilities. Let $\mathbb{Q}_N$ denote the one-period risk-neutral measure associated with the SDF M. By definition

$$\frac{d\mathbb{Q}_N}{d\mathbb{P}} = \frac{M}{\mathbb{E}[M]}$$

Combining (18) with (10) yields

$$\begin{aligned} p_{\mathbb{Q}_N} =& \mathbb{Q}_N(J = j) = \frac{\mathbb{E}\left[\mathbb{1}_{\{J=j\}} \frac{d\mathbb{Q}_N}{d\mathbb{P}}\right]}{\mathbb{E}\left[\frac{d\mathbb{Q}_N}{d\mathbb{P}}\right]} = \frac{\mathbb{E}\left[\mathbb{1}_{\{J=j\}} e^{-(\gamma+\theta)(z+J)}\right]}{\mathbb{E}\left[e^{-(\gamma+\theta)(z+J)}\right]} \\ =& \frac{p\, e^{-(\gamma+\theta) j}}{1 - p + p\, e^{-(\gamma+\theta) j}} \end{aligned} \tag{30}$$

where the last line follows from the fact that z and J are independent. This yields the risk-neutral disaster-state probability (or jump mass). Intuitively it measures how high the market treats the probability of an event occurring relative to its physical ("true") probability. That is, $p_{\mathbb{Q}_N} > p$ implies the market treats the probability of a disaster occurring higher as it is due to tail aversion.[4] This constitutes a market-side perspective of how strong tail aversion is. As $p_{\mathbb{Q}_N}$ overweights rare losses relative to the

[4]Note that disasters are over-weighted under $\mathbb{Q}_N$ whenever $\gamma + \theta > 0$ and more so the more negative j is.

physical measure this endogenous reweighting parallels the probability weighting in cumulative prospect theory (Tversky and Kahneman, 1992). It is also closely related in spirit to the prospect-theory asset-pricing mechanisms in Barberis, Huang, and Santos (2001). However, in the given context it arises from an entropic change of measure in an expected-utility setting.

4.3 Intuition

To recap, prices satisfy $1 = \mathbb{E}_t\left[M_{t+1}R_{t+1}\right]$ with

$$M_{t+1} = \frac{e^{-\delta-(\gamma+\theta)g_{t+1}}}{\mathbb{E}_t\left[e^{-\theta g_{t+1}}\right]}$$

where $g_{t+1} = \ln\frac{C_{t+1}}{C_t}$ is consumption growth, γ is relative risk aversion, and $\theta > 0$ implements an entropic tilt that overweights bad consumption states ("left tails"). Intuitively, θ acts like an additional sensitivity of marginal utility to g_{t+1}. More precisely, the tilt increases the local loading of $\ln M_{t+1}$ on g_{t+1} from $-\gamma$ to $-(\gamma+\theta)$ without altering period utility curvature. Consumption growth itself has a Gaussian diffusive part $z_{t+1} \overset{i.i.d.}{\sim} \mathcal{N}\left(\mu,\sigma^2\right)$ and a rare-disaster jump $J_{t+1} \in \{0; j\}$ with $j < 0$ occurring with probability p. Log asset returns co-move with the diffusive part of consumption growth with coefficient ρ and co-jump by $qj < 0$ in disasters. Thus two channels create negative $\mathrm{Cov}\left[M_{t+1}, R_{t+1}\right]$

1. **Diffusive channel:** when z_{t+1} is low, M_{t+1} is high and (for $\rho > 0$) R_{t+1} is low, increasing risk compensation.

2. **Jump channel:** in rare disasters C_{t+1} and prices drop. If valuation ratios compress on impact, the co-jump qj magnifies equity losses beyond the one-period dividend shortfall ($q > 1$), so payoffs are particularly poor exactly when M_{t+1} is large.

Both channels can be seen immediately in the log risk-free rate, given by

$$\ln R_f = \delta + \gamma\mu - \frac{1}{2}\gamma^2\sigma^2 \quad \underbrace{-\gamma\,\theta\,\sigma^2}_{\text{diffusion interaction}} \quad \underbrace{+\ln\left[\frac{1-p+p\,e^{-\theta j}}{1-p+p\,e^{-(\gamma+\theta)j}}\right]}_{\text{jump ratio}}$$

The first three terms are the standard CRRA/lognormal elements of the consumption-based asset pricing model. The two new ingredients lower $\ln R_f$. First, the cross term $-\gamma\theta\sigma^2$ lowers the log risk-free rate because the tilt and curvature interact under diffusion. Second, the jump ratio lowers the log risk-free rate as for $\theta > 0$ and $j < 0$ the denominator is larger than the numerator. Hence the model naturally produces a low risk-free rate when $p > 0$ and $\theta > 0$, echoing the risk-free rate puzzle. Intuitively, a high tail aversion θ implies investors preferring the "safe haven" of the risk-free rate over risky assets.

For the risky asset, the small-risk approximation again highlights both channels

$$\mathrm{SR} \approx \left[\underbrace{(\gamma+\theta)\, \sigma \sigma_r \rho}_{\text{diffusive covariance}} \underbrace{- p \left(e^{qj} - 1\right) \left(e^{-(\gamma+\theta)j} - 1\right)}_{\text{jump covariance}} \right] \left[\sqrt{\sigma_r^2 + p\,(1-p)\,\left(e^{qj}-1\right)^2} \right]^{-1}$$

The diffusion part scales with $(\gamma+\theta)$ and the exposure ρ. The jump covariance term also raises the Sharpe ratio as $\left(e^{qj}-1\right)\left(e^{-(\gamma+\theta)j}-1\right) < 0$.[5] Subsequently, the jump part is positive whenever $qj < 0$ and $\gamma+\theta > 0$, because returns fall in disasters while the SDF spikes. Thus either higher θ (more tail aversion) or larger q (stronger valuation compression on impact) raises required premia and Sharpe ratios. This can most clearly be observed in the HJ bound

$$\mathrm{SR}_{\max} = \sqrt{\frac{\mathbb{E}\left[M_{t+1}^2\right]}{\mathbb{E}\left[M_{t+1}\right]^2} - 1} = \sqrt{e^{(\gamma+\theta)^2\sigma^2}\, \frac{1-p+p\,e^{-2(\gamma+\theta)j}}{\left(1-p+p\,e^{-(\gamma+\theta)j}\right)^2} - 1}$$

where both the tilt θ and rare disasters ($p > 0$, $j < 0$) increase the volatility of M_{t+1} relative to its mean, thus pushing out the HJ bound and helping reconcile observed Sharpe ratios while maintaining modest γ.

5 Bounds on risk aversion in the presence of tail aversion

The model also implies natural bounds on risk aversion given non-negative values for the tail aversion parameter θ and at least proportional valuation compression q. The following Lemma establishes an upper bound on risk aversion implied by the risk-free rate.

Lemma 3. *Let $\ln R_f < \delta$ and fix parameters (μ, σ, p, j). Then the risk-free rate implies the existence and uniqueness of an upper bound on risk aversion $\gamma^\dagger$. The maximum admissible value is attained when tail aversion is zero, i.e. $\theta = 0$.*

Proof. See Appendix B.1.

□

The intuition behind this result is straightforward. Increasing either risk aversion or tail aversion always lowers the risk-free rate. Hence the largest level of risk aversion consistent with the observed risk-free rate occurs when tail aversion is shut down. The

[5] To see this, note that $j < 0$, $q > 1$, and thus $qj < 0$, $e^{qj} - 1 < 0$, and $e^{-(\gamma+\theta)j} - 1 > 0$.

following Proposition establishes the functional relationship between tail aversion and risk aversion.

Proposition 4. *Let* $\ln R_f < \delta$ *and fix parameters* (μ, σ, p, j)*. Then for each* $\gamma \in (0, \gamma^{\dagger})$ *there exists a unique* $\theta(\gamma) > 0$ *solving the risk-free rate equation* (20) *with* $\theta'(\gamma) < 0$ *and*

$$\theta'(\gamma) \sim -\frac{\delta - \ln R_f}{\gamma^2 \sigma^2} \qquad \text{for } \gamma \downarrow 0$$

Proof. See Appendix B.2.

□

That is, in order to match the same risk-free rate, a reduction in risk aversion must be offset by stronger tail aversion. Conversely, higher risk aversion requires less tail aversion. Thus the mapping from risk aversion to tail aversion is downward sloping. Taken together, the risk-free rate pins down a mapping of $\theta(\gamma) > 0$ for a given set of parameters within the domain $\gamma \in (0, \gamma^{\dagger})$. To obtain the bounds on γ define $A(q) \equiv 1 - p + p\, e^{qj}$, $B(q) \equiv 1 - p + p\, e^{2qj}$, $m \equiv \mathbb{E}[R]$, and $s^2 \equiv \mathrm{Var}[R]$. Applying these definitions and re-arranging (22) and (23) yields

$$\sigma_r(q) = \sqrt{\ln\left(1 + \frac{s^2}{m^2}\right) + 2\ln A(q) - \ln B(q)} \quad \text{and} \quad \mu_r + \frac{1}{2}\sigma_r^2 = \ln m - \ln A(q)$$

The following Lemma ensures positive volatility of the diffusive part of asset returns.

Lemma 4. *A sufficient condition that ensures* $\sigma_r > 0$ *for all* $q \geq 1$ *is given by*

$$\frac{s}{m} > \sqrt{\frac{p}{1-p}} \tag{31}$$

Proof. See Appendix B.3.

□

The inequality ensures that, even if jump risk absorbs a lot of the overall volatility, the lognormal diffusion still carries strictly positive volatility. Intuitively, the observed equity volatility cannot be "too low" relative to its mean if the diffusion part is to remain active. Substituting the results into the asset pricing equation, given by (21), collecting terms, and applying Proposition 4 yields

$$
\begin{aligned}
F(q,\gamma) \equiv & \delta + \gamma\mu - \frac{1}{2}\gamma^2\sigma^2 - \gamma\,\theta(\gamma)\,\sigma^2 \\
& + (\gamma + \theta(\gamma))\,\rho\,\sigma\sqrt{\ln\left(1+\frac{s^2}{m^2}\right) + 2\ln A(q) - \ln B(q)} \\
& + \ln\left[1 - p + p\,e^{-\theta(\gamma)\,j}\right] - \ln\left[1 - p + p\,e^{-(\gamma+\theta(\gamma)-q)j}\right] \\
& - \ln m + \ln A(q) = 0
\end{aligned}
\tag{32}
$$

Feasibility requires jointly solving (20) and (32) for γ and $\theta(\gamma)$ over the admissible range of values $\gamma \in (0,\ \gamma^\dagger)$, conditional on $q \geq 1$ and $\sigma_r > 0$. The following assumption ensures non-degeneracy

Assumption 1. *Let $\theta(\gamma)$ be the unique mapping solving the risk-free equation (20) for $\gamma \in (0,\ \gamma^\dagger)$, and define the feasible set*

$$
\Gamma \equiv \left\{\gamma \in \left(0,\ \gamma^\dagger\right) : \exists\, q \geq 1 \text{ with } F(q,\gamma) = 0\right\}
$$

On any connected component of the branch

$$
\mathcal{B} \equiv \{(\gamma, q) \in \Gamma \times [1, \infty) : F(q,\gamma) = 0\}
$$

I assume non-degeneracy in the partial derivative, that is

$$
\frac{\partial F(q,\gamma)}{\partial q} \neq 0 \quad \text{for all } \ (\gamma, q) \in \mathcal{B}
$$

This regularity condition rules out knife-edge cases in which the co-jump parameter is locally indeterminate. On any feasible branch, the asset-pricing equation then selects a unique co-jump level for a given risk aversion. The following Proposition establishes the functional relationship between co-jump and risk aversion.

Proposition 5. *Under the restriction implied by Lemma 4, $\mu > -p\,j$, and $\rho > 0$ let $\theta(\gamma)$ be given by Proposition 4 and let $q(\gamma) \geq 1$ solve (32) wherever a solution exists. Then on any connected component of the feasible set $\Gamma := \left\{\gamma \in \left(0,\ \gamma^\dagger\right) : \exists\, q \geq 1 \text{ with } F(q,\gamma) = 0\right\}$ the map $\gamma \mapsto q(\gamma)$ is C^1 and single-peaked. That is, there exists a unique $\gamma^\star \in \left(0,\ \gamma^\dagger\right)$ such that $q'(\gamma^\star) = 0$, and $q'(\gamma)$ switches sign once on each component.*

Proof. See Appendix B.4. □

Intuitively, holding the equity mean and variance fixed, a larger co-jump q shifts risk from the smooth diffusion to rare jumps. To keep the same Sharpe-relevant moments, $\sigma_r(q)$ must fall with q, so the diffusion price-of-risk component $(\gamma + \theta)\,\rho\,\sigma\,\sigma_r(q)$ declines as q rises, given $\rho > 0$. Since the risk-free rate is held fixed, Proposition 4 implies that tail-aversion $\theta(\gamma)$ is decreasing in γ. At very low γ, the required θ to match a given risk-free rate is high, so the SDF's diffusion loading $\omega(\gamma) \equiv \gamma + \theta(\gamma)$ is already large. Additionally, disasters are heavily priced via the large $\theta(\gamma)$. Consequently, little additional valuation compression is needed and q sits near its lower bound. As γ increases from this region, $\theta(\gamma)$ falls sharply in accordance with Proposition 4, so $\omega(\gamma)$ initially declines. To keep the same equity mean and variance, the model compensates by shifting risk into jumps, which raises the required q. Once $\theta(\gamma)$ has fallen enough, $\omega(\gamma)$ begins to rise with γ and the diffusion channel regains strength, so the needed co-jump falls again. This interaction, i.e. high θ at low γ and then strengthening diffusion at higher γ, generates a non-monotonic relation between γ and q with a single interior maximum for $q(\gamma)$. Finally, the results for the admissible bounds on risk aversion are summarized in the following Proposition.

Proposition 6. *For given moments (R_f, m, s) with $\ln R_f < \delta$, $m > 0$, $s > 0$, $\frac{s}{m} > \sqrt{\frac{p}{1-p}}$, $\mu > -pj$, and $\rho > 0$ the system consisting of (20) and (32) admits a solution $(\theta(\gamma) \geq 0, q(\gamma) \geq 1)$ if and only if $\gamma \in (\gamma_{\min}, \gamma_{\max})$, where the endpoints are unique and characterized as follows:*

1. *$\gamma_{\min}$ is the unique smallest root of $F(1,\gamma) = 0$ with $\theta(\gamma)$ given by solution to (20)*
2. *$\gamma_{\max} = \min\{\gamma^{\dagger}, \gamma^{+}\}$, where*
   - *i) $\gamma^{\dagger}$ solves (20) for $\theta = 0$*
   - *ii) γ^{+} is the unique largest root of $F(1,\gamma) = 0$ with $\theta(\gamma)$ given by the solution to (20)*

Proof. See Appendix B.5. □

Proposition 6 pins down a strict admissible band for risk aversion. A feasible pair $(\theta(\gamma) \geq 0, q(\gamma) \geq 1)$ exists if and only if γ lies between a lower root where valuation compression just vanishes (i.e., $q = 1$) and an upper bound determined either by the upper root where valuation compression just vanishes or the ceiling $\gamma^{\dagger}$ implied by Lemma 3, whichever is smaller. Economically, because the risk-free rate condition forces tail aversion $\theta(\gamma)$ to fall with γ, the compression needed to match equity moments first rises above one and then falls back to one, so feasibility fails outside $(\gamma_{\min}, \gamma_{\max})$.

Table 1: exogenous parameters

Parameter	Value	Target
δ	0.03	discount rate (implies annual $\beta = e^{-\delta} \approx 0.97$)
γ	3	risk aversion
ρ	0.2	diffusion correlation consumption growth and returns
μ	0.025	consumption growth p.a. outside of disasters
σ	0.02	consumption growth volatility p.a. outside of disasters
p	0.017	disaster probability p.a.
j	$\ln(1-0.29)$	consumption drop in disaster

6 Quantitative analysis

The following section presents a quantitative evaluation of the model. After disciplining the model to rare disaster data by Barro (2006) an indirect inference strategy backs out the size of the tail-aversion parameter θ along with the size of the valuation compression q and the diffusion volatility of returns σ_r. Remaining preference and diffusion parameters are set to long-run benchmarks and kept within conventional ranges. As the model also nests standard cases, the analysis proceeds via neutralizing parts of the newly introduced modifications via either setting θ or p to zero thus obtaining information on the importance of the model's channels in isolation. Last, I perform sensitivity checks to show that the results are robust to changing the diffusive correlation and the degree of risk aversion.

6.1 Calibration and indirect inference

I calibrate the model to long-run moments from Jordà, Knoll, Kuvshinov, Schularick, and Taylor (2019), who construct a long-run series of returns for 16 developed countries and construct averages weighted by GDP. The unit of observation is one year. The values of δ, γ, and ρ are taken directly from standard values of the literature. The parameters of consumption growth, namely the diffusive mean and volatility (μ and σ) as well as the probability of disaster occurrence and the associated average drop in consumption are taken from Barro (2006) and set to 2.5%, 2%, 1.7%, and 29%, respectively. Table 1 summarizes the values.

Equations (18)–(23) imply a convenient inversion to match return mean/variance while keeping (p, j) fixed. To obtain a value for θ I solve (20) for θ given the calibrated parameter values targeting a real net risk-free rate of 1.3%. Defining the gross mean $m \equiv \mathbb{E}[R]$ with net return $m - 1$ and variance $s^2 \equiv \mathrm{Var}[R]$ re-arranging (22) and (23) yields

Table 2: baseline results

Parameter/Statistic	Value	Data/Targets
θ	3.267	–
σ_r	0.2	–
Implied equity loss	44%	–
Net risk-free rate	0.013*	0.013
Net asset return	0.071*	0.071
Asset volatility	0.22*	0.22
Sharpe ratio	0.264	0.264
HJ bound	0.881	$\geq$ 0.264
D_{KL}	0.024	–
TV	11 pp	–
$p_{\mathbb{Q}_N}$	13%	–

$$\sigma_r(q) = \sqrt{\ln\left(1+\frac{s^2}{m^2}\right) + 2\ln A(q) - \ln B(q)} \quad \text{and} \quad \mu_r(q) + \frac{1}{2}\sigma_r(q)^2 = \ln m - \ln A(q)$$

where again $A(q) \equiv 1 - p + p\, e^{qj}$ and $B(q) \equiv 1 - p + p\, e^{2qj}$ while $\mu_r(q) + \frac{1}{2}\sigma_r(q)^2$ is given by (21) as

$$\begin{aligned} \mu_r(q) + \frac{1}{2}\sigma_r(q)^2 =& \delta + \gamma\mu - \frac{1}{2}\gamma^2\sigma^2 - \gamma\,\theta\,\sigma^2 + (\gamma+\theta)\,\rho\,\sigma\,\sigma_r(q) \\ &+ \ln\left[\frac{1-p+p\,e^{-\theta\, j}}{1-p+p\,e^{-(\gamma+\theta-q)j}}\right] \end{aligned}$$

The values of q, σ_r, and μ_r are then obtained via solving this system of equations for asset market moments targeting a real net equity return of 7.1% with 22% volatility.

6.2 Results and discussion

Table 2 summarizes the core results. The chosen parameterization matches the market return and volatility as well as the risk-free rate by construction. As a result, the Sharpe ratio implied by the model is identical to the one provided by the data. Note that the one-period Kullback-Leibler divergence implied by the model is only 0.024 nats. That is, from an information perspective the tail aversion is almost not detectable and the model is informationally virtually indistinguishable from the non-tail adjusted (CRRA) case. However, from a market perspective the tail aversion is substantial as $p_{\mathbb{Q}_N} \approx 13\%$ and thus almost an order of magnitude larger than the physical probability $p = 1.7\%$.

To achieve this for the given parameterization the model requires the following quantitative features

1. The tail adjustment parameter has to be similar in size to baseline risk aversion.
2. The implied equity loss on impact from a disaster has to be around 44%.
3. The contribution of the diffusion part to overall asset volatility is around 92% as

$$\frac{\mathbb{E}\left[\mathrm{Var}\left[R|J\right]\right]}{\mathrm{Var}\left[R\right]} = \frac{B(q)\left(e^{\sigma_r^2}-1\right)}{B(q)e^{\sigma_r^2}-A(q)^2} \approx 92\%$$

Taken together this implies that a model with tail adjustment in the SDF and the potential for extreme events is able to quantitatively reconcile both the equity premium puzzle and the risk-free rate puzzle without the need to resort to unreasonably high degrees of risk aversion. Note that in addition to the low D_{KL} even the combined local log-SDF loading on the diffusion part, i.e. $\gamma + \theta \approx 6.27$, does not imply a very high value of risk aversion. Using (13) allows evaluating the size of losses implied by the co-jump parameter q as $L = 1 - e^{qj}$. Historical episodes suggest j between 15% and $\approx$ 60% in major crises (Barro and Ursúa, 2008a; Nakamura, Steinsson, Barro, and Ursúa, 2013), while equity drawdowns on impact range widely. Peak-to-trough losses of 40% - 60% are common in major crashes. Hence, $L \approx 44\%$ is well within values consistent with large crisis episodes.

6.3 Channels

In order to explore the channels highlighted in Section 4 I run additional experiments that neutralize parts of the model. As a first pass, I re-run the model without tail adjustment or extreme events (i.e. $\theta = p = 0$) to obtain the standard benchmark of the consumption-based asset pricing model. Table 3 presents the results of this experiment in column 1. The parameter σ_r is calibrated to match asset volatility, which is matched by construction. The remaining moments in column 1 highlight the well-known problems of the standard consumption-based asset pricing model. The risk free-rate is an order of magnitude too high for standard parameters and the Sharpe ratio is missed by more than an order of magnitude. As there are neither tail events nor tail adjustment all tail aversion measures are zero.

As shown in column 2, allowing for a tail adjustment in the SDF that is calibrated to match the risk-free rate fixes most of these issues. However, the model has to resort to an unreasonably high degree of adjustment ($\theta \approx 75$) to achieve this. This result is not surprising as in the absence of truly extreme events the aversion to left-tail realizations has to be profoundly high in order to justify a net risk-free rate of 1.3% in real terms. Naturally, this causes the one-period Kullback-Leibler divergence to grow substantially such that tail aversion is now easily detectable while $p_{\mathbb{Q}_N}$ does not move as this specification contains no disaster events.

Table 3: importance of channels

Parameter/Statistic	(1) Standard model	(2) $p = 0$	(3) $\theta = 0$	(4) Full	Data/Targets
θ	0	75.236	0	3.267	–
p	0	0	5.27%	1.7%	–
σ_r	0.2	0.2	0.18	0.2	–
Implied equity loss	0	0	44%	44%	–
Net risk-free rate	0.109	0.013*	0.013*	0.013*	0.013
Net asset return	0.111	0.079	0.054	0.071*	0.071
Asset volatility	0.22*	0.22*	0.22*	0.22*	0.22
Sharpe ratio	0.012	0.3	0.186	0.264	0.264
HJ bound	0.06	3.251	0.372	0.881	≥ 0.264
D_{KL}	0	1.132	0	0.024	–
TV	0 pp	75 pp	0 pp	11 pp	–
$p_{\mathbb{Q}_N}$	0%	0%	13%	13%	–

* denotes statistic that the parameters were chosen to replicate

Conversely, column 3 shows that allowing for disasters, whose probability is again calibrated to match the risk-free rate, while abstracting from tail adjustment in the SDF and retaining the same degree of asset losses as in the baseline results, also improves the Sharpe ratio considerably. However, as can be inferred from column 3 this comes at the cost of a very high probability of disaster ($p \approx 5.3\%$) while generating lower volatility of the diffusion process ($\sigma_r \approx 0.18\%$). Again, these results are not surprising as a high probability of disaster is required to suppress the net risk-free rate to 1.3%. Subsequently, as the disaster probability is now higher, more of asset volatility will stem from the disaster jumps occurring with a higher probability instead of from the diffusion process. As with column (1) there is no tail adjustment and thus all information criteria detecting tail aversion are zero. However, $p_{\mathbb{Q}_N}$ increases substantially due to the presence of disasters.

Finally, column 4 shows that the results of a low risk-free rate and a high Sharpe ratio can be reconciled by allowing a moderate tail adjustment paired with a low probability of disaster that generates large asset losses.

6.4 Sensitivity

The quantitative results are robust to variation in ρ, i.e. the correlation between the diffusion parts of consumption growth and asset returns. As the risk-free rate is independent of ρ variation in the parameter does not alter it and hence also does not influence the size of θ. Conversely, (21) shows that equity premia increase with $(\gamma + \theta)\, \rho\, \sigma\, \sigma_r$. Consequently, lowering ρ forces more weight on disaster co-jumps,

Table 4: sensitivity

Parameter/Statistic	(1) Benchmark	(2) $\rho = 0.05$	(3) $\rho = 0.8$	(4) $\gamma = 1$	(5) $\gamma = 7$
θ	3.267	3.267	3.267	5.405	0.461
σ_r	0.2	0.19	0.2	0.2	0.2
Implied equity loss	44%	47%	31%	42%	29%
D_{KL}	0.024	0.024	0.024	0.102	0.0003
TV	11 pp	11 pp	11 pp	23 pp	1 pp
$p_{\mathbb{Q}_N}$	13%	13%	13%	13%	18%

raising q and making equity losses more front-loaded in bad states as can be seen in column 2 of Table 4. On the other hand, raising ρ shifts premia to the diffusion channel and reduces the q needed to match m and s (column 3 in Table 4). In order to ensure that the shift to diffusion does not give rise to $q \leq 1$ in the indirect inference approach the following additional restriction is needed

$$\rho < \rho_{\max} = \frac{\ln m - \ln A(1) - \left(\delta + \gamma\mu - \frac{1}{2}\gamma^2\sigma^2 - \gamma\,\theta\,\sigma^2\right) - \ln\left[\frac{1-p+p\,e^{-\theta j}}{1-p+p\,e^{-(\gamma+\theta-1)j}}\right]}{(\gamma+\theta)\;\sigma\,\sqrt{S(1)}} \tag{33}$$

where θ is obtained by solving for the risk-free rate, which is independent of ρ. For the baseline parameterization this yields $\rho_{\max} \approx 0.9$, which is larger than any empirical value typically associated with this correlation. Within the admissible range of $\rho \in (0, \rho_{\max})$ the quantitative importance of this channel is muted as indicated by the results in columns 2 and 3 in Table 4. That is, the results do not strongly depend on any specific size of the correlation between the diffusion parts of consumption growth and asset returns.

The results are also robust to variation in risk aversion γ within the admissible bounds derived in Section 5. Under the baseline parameters solving (32) conditional on (20) yields $\gamma_{\max} \approx 7.08$ and $\gamma_{\min} \approx 0.46$, which are in line with standard conjectures about the upper and lower bound of risk aversion. Taken together the admissible range of risk aversion is given by $\gamma \in (\gamma_{\min} \approx 0.46,\ \gamma_{\max} \approx 7.08)$. Within those values (θ, q, σ_r) exist and match (R_f, m, s). As derived in Proposition 4, the expression for the risk-free rate implies that, holding R_f fixed, $\theta\,(\gamma)$ is decreasing in γ for $j < 0$ and $p > 0$ in the empirically relevant range. This pattern is shown by columns 4 and 5 in Table 4. For $\gamma = 1$ the calibration requires $\theta \approx 5.4$, while for $\gamma = 7$ it requires $\theta \approx 0.46$ (both with unchanged $\sigma_r \approx 0.2$). At the same time, the asset pricing equation (32) shows that a rise in γ increases the diffusion loading $(\gamma+\theta)\,\rho\,\sigma\,\sigma_r$ but simultaneously lowers θ, so the net effect on $(\gamma+\theta)$ is muted in practice. To keep the equity mean and variance

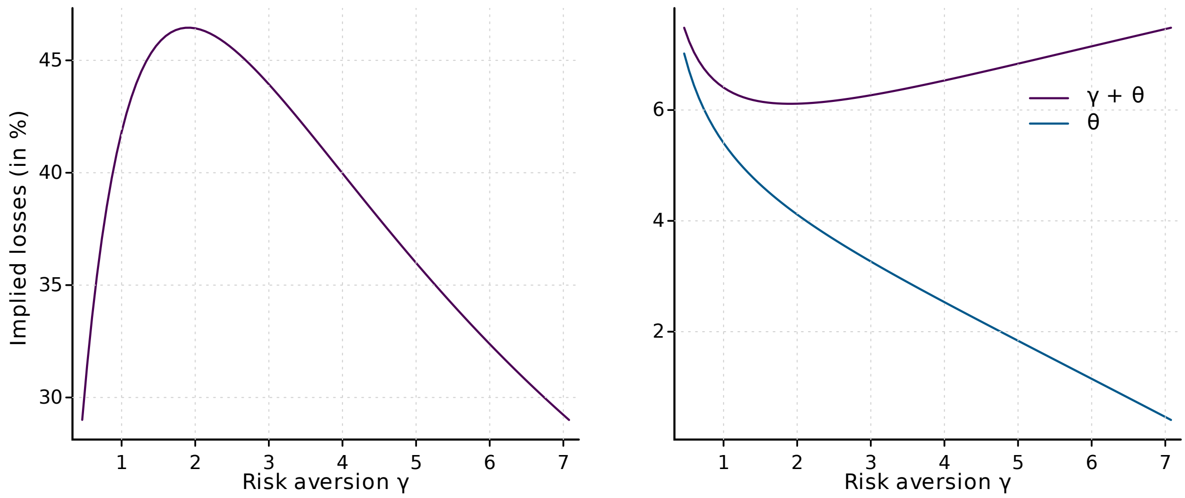

(a) Implied losses $L = 1 - e^{qj}$

(b) Implied θ

Figure 1: Reaction of valuation compression and tail aversion to risk aversion

fixed, the system rebalances by adjusting q. Columns 4 and 5 of Table 4 as well as panel (a) of Figure 1 illustrate the non-monotonic behavior of the adjustment of q laid out in Proposition 5.

A large microeconomic literature finds values of relative risk aversion at or below one in laboratory and field settings. For instance, Holt and Laury (2002) estimate CRRA well below unity at moderate stakes (with pronounced incentive effects) while the calibration critique of Rabin (2000) implies that small-stakes curvature must be modest if one wishes to avoid counterfactual large-stakes behavior. Surveying field data, Barseghyan, Molinari, O'Donoghue, and Teitelbaum (2018) conclude that moderate risk aversion and substantial heterogeneity are the norm across insurance and portfolio contexts. At the other end of the spectrum, macroeconomic evidence sometimes implies large curvature when CRRA also proxies the intertemporal elasticity of substitution. As an example, Best, Cloyne, Ilzetzki, and Kleven (2020) estimate an average intertemporal elasticity of substitution near 0.1 using quasi-experimental "mortgage notches", which under CRRA maps to $\gamma \approx 10$. The bounds-based exercise here is qualitatively consistent with both possibilities. As can be observed from comparing panel (a) and panel (b) of Figure 1 the implied valuation compression (and hence implied loss) is highest at intermediate ranges of risk aversion. A higher γ permits a smaller q (less valuation compression) while maintaining the same (m, s) as high risk aversion suffices to match the moments of interest. On the other hand, a low γ is also consistent with less valuation compression as higher tail aversion θ depresses the risk-free rate and raises the equity return. Simultaneously, raising θ implies a stronger penalty on the one-period Kullback-Leibler divergence while lowering reduces it. Thus strong tail aversion combined with low risk aversion is more easily detectable than moderate/high risk aversion paired with moderate/low tail aversion. Taken together, the quantitative fit is robust to a large range of admissible values of γ and admits fitting the data even for very low levels of risk aversion.

7 Concluding remarks

This paper offers a compact asset-pricing model that unifies two empirically grounded mechanisms: rare consumption disasters and an entropic, risk-sensitive tilt in the pricing kernel. Analytically, the SDF's diffusion loading $\gamma + \theta$ and the disaster co-jump parameter q deliver closed-form expressions for the risk-free rate, the Sharpe ratio, and the HJ bound. The model explains why premia are high exactly when the economy is weak. The tilt raises the price of diffusion risk in bad times, and disaster co-jumps depress payoffs precisely when marginal utility is high.

The quantitative exercise shows that a small probability of sizable disasters, paired with moderate tail aversion θ, can match equity premia and the low real risk-free rate without extreme curvature. In fact, from an information perspective the model is almost indistinguishable from CRRA preferences. The implied equity loss on impact, $1 - e^{qj}$, is in line with large historical crashes, and the decomposition of premia into diffusion and jump components clarifies the role of correlation and valuation compression at disaster times.

There are natural next steps. On the empirical side, the magnitudes of (p, j, q) can be further disciplined with cross-country disaster data and event-based measures of valuation compression. On the theoretical side, embedding the one-period tilt into a model with endogenous investment and intermediaries would allow joint study of the term structure, credit spreads, and cross-sectional asset pricing. Moreover, the analysis intentionally models disasters as one-period events. This abstraction preserves closed-form mixture-lognormal pricing for key moments and keeps the SDF transparent. This comes with the cost of not being able to account for crisis dynamics, term-structure patterns during slumps, or volatility clustering. A natural next step here is to allow time variation in pricing states. Allowing partial recoveries and time-varying disaster intensities (e.g., Nakamura, Steinsson, Barro, and Ursúa, 2013), variable disaster risk p_t in the spirit of Gabaix (2012), or a time-varying tail aversion θ_t that tightens in bad times would generate countercyclical prices of risk, return predictability, and more realistic volatility dynamics, while preserving the parsimony of the one-period tilt and the closed-form conditional moments developed here. Finally, formal estimation could exploit the analytic inversion mapping moments to structural parameters.

Overall, the results support the simple message that accounting for tail risk and for state-dependent pricing weights suffices to reconcile first-order asset-pricing moments while preserving parsimony and interpretability.

References

Aiyagari, S. R. (1994): “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics*, 109(3), 659–684.

Ang, A., R. J. Hodrick, Y. Xing, and X. Zhang (2006): “The Cross-Section of Volatility and Expected Returns,” *Journal of Finance*, 61(1), 259–299.

Bansal, R., D. Kiku, and A. Yaron (2012): “An Empirical Evaluation of the Long-Run Risks Model for Asset Prices,” *Journal of Finance*, 67(6), 1915–1963.

Bansal, R., and A. Yaron (2004): “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *Journal of Finance*, 59(4), 1481–1509.

Barberis, N., M. Huang, and T. Santos (2001): “Prospect Theory and Asset Prices,” *The Quarterly Journal of Economics*, 116(1), 1–53.

Barro, R. J. (2006): “Rare Disasters and Asset Markets in the Twentieth Century,” *Quarterly Journal of Economics*, 121(3), 823–866.

——— (2009): “Rare Disasters, Asset Prices, and Welfare Costs,” *American Economic Review*, 99(1), 243–264.

Barro, R. J., and J. F. Ursúa (2008a): “Consumption Disasters in the Twentieth Century,” *American Economic Review*, 98(2), 58–63.

——— (2008b): “Macroeconomic Crises since 1870,” *Brookings Papers on Economic Activity*, 2008, 255–335.

Barseghyan, L., F. Molinari, T. O’Donoghue, and J. C. Teitelbaum (2018): “Estimating Risk Preferences in the Field,” *Journal of Economic Literature*, 56(2), 501–564.

Benartzi, S., and R. H. Thaler (1995): “Myopic Loss Aversion and the Equity Premium Puzzle,” *Quarterly Journal of Economics*, 110(1), 73–92.

Best, M. C., J. S. Cloyne, E. Ilzetzki, and H. J. Kleven (2020): “Estimating the Elasticity of Intertemporal Substitution Using Mortgage Notches,” *Review of Economic Studies*, 87(2), 656–690.

Brunnermeier, M. K., and Y. Sannikov (2014): “A Macroeconomic Model with a Financial Sector,” *American Economic Review*, 104(2), 379–421.

Campbell, J. Y., and J. H. Cochrane (1999): “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, 107(2), 205–251.

Campbell, J. Y., M. Lettau, B. G. Malkiel, and Y. Xu (2001): "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk," *Journal of Finance*, 56(1), 1–43.

Chetty, R. (2006): "A New Method of Estimating Risk Aversion," *American Economic Review*, 96(5), 1821–1834.

Chiappori, P.-A., and M. Paiella (2011): "Relative Risk Aversion Is Constant: Evidence from Panel Data," *Journal of the European Economic Association*, 9(6), 1021–1052.

Cochrane, J. H. (2005): *Asset Pricing: Revised Edition*. Princeton University Press, Princeton, NJ.

——— (2017): "Macro-Finance," *Review of Finance*, 21(3), 945–985.

Cohen, A., and L. Einav (2007): "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, 97(3), 745–788.

Constantinides, G. M., and D. Duffie (1996): "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 104(2), 219–240.

Epstein, L. G., and S. E. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57(4), 937–969.

——— (1991): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis," *Journal of Political Economy*, 99(2), 263–286.

Gabaix, X. (2012): "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," *Quarterly Journal of Economics*, 127(2), 645–700.

Hansen, L. P., and R. Jagannathan (1991): "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99(2), 225–262.

Hansen, L. P., and T. J. Sargent (2001): "Robust Control and Model Uncertainty," *American Economic Review*, 91(2), 60–66.

——— (2008): *Robustness*. Princeton University Press, Princeton, NJ.

He, Z., and A. Krishnamurthy (2013): "Intermediary Asset Pricing," *American Economic Review*, 103(2), 732–770.

Heaton, J., and D. Lucas (1996): "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing," *Journal of Political Economy*, 104(3), 443–487.

Holt, C. A., and S. K. Laury (2002): "Risk Aversion and Incentive Effects," *American Economic Review*, 92(5), 1644–1655.

Jordà, Ò., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor (2019): "The Rate of Return on Everything, 1870–2015," *The Quarterly Journal of Economics*, 134(3), 1225–1298.

Kahneman, D., and A. Tversky (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47(2), 263–291.

Kullback, S., and R. A. Leibler (1951): "On Information and Sufficiency," *The Annals of Mathematical Statistics*, 22(1), 79–86.

Lettau, M., and S. Ludvigson (2001): "Consumption, Aggregate Wealth, and Expected Stock Returns," *Journal of Finance*, 56(3), 815–849.

Lucas Jr., R. E. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46(6), 1429–1445.

Maenhout, P. J. (2004): "Robust Portfolio Rules and Asset Pricing," *Review of Financial Studies*, 17(4), 951–983.

Martin, I. W. R. (2008): "Disasters and the Welfare Cost of Uncertainty," *American Economic Review*, 98(2), 74–78.

——— (2013): "Consumption-Based Asset Pricing with Higher Cumulants," *Review of Economic Studies*, 80(2), 745–773.

Mehra, R., and E. C. Prescott (1985): "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15(2), 145–161.

——— (2003): "The Equity Premium in Retrospect," *Journal of Economic Perspectives*, 17(1), 239–258.

Nakamura, E., J. Steinsson, R. Barro, and J. Ursúa (2013): "Crises and Recoveries in an Empirical Model of Consumption Disasters," *American Economic Journal: Macroeconomics*, 5(3), 35–74.

Rabin, M. (2000): "Risk Aversion and Expected-Utility Theory: Calibration Theorems," *Econometrica*, 68(5), 1281–1292.

Rietz, T. A. (1988): "The Equity Risk Premium: A Solution," *Journal of Monetary Economics*, 22(1), 117–131.

Storesletten, K., C. I. Telmer, and A. Yaron (2004): "Consumption and Risk Sharing Over the Life Cycle," *Journal of Political Economy*, 112(4), 722–745.

Tversky, A., and D. Kahneman (1992): "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5(4), 297–323.

Wachter, J. A. (2013): "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?," *Journal of Finance*, 68(3), 987–1035.

Weil, P. (1989): "The Equity Premium Puzzle and the Risk-Free Rate Puzzle," *Journal of Monetary Economics*, 24(3), 401–421, NBER Working Paper No. 2829.

A A simple model of valuation compression

Let $\kappa_t \equiv \frac{P_t}{D_t}$ denote the price-dividend ratio and $\Delta d_{t+1} \equiv \ln D_{t+1} - \ln D_t$ dividend-growth. Then the standard Campbell-Shiller log-linearization implies

$$\ln R_{t+1} \approx k + \eta \ln \kappa_{t+1} - \ln \kappa_t + \Delta d_{t+1} \tag{A.1}$$

where k and $\eta \in (0,1)$ are log-linearization constants determined by the unconditional mean of $\ln \kappa_t$. For the reduced form disaster mapping, a consumption (and dividend) contraction of fractional size $x \in (0,1)$ is given by $j = \ln(1-x) < 0$. If equity loses a fraction $L \in (0,1)$ on impact, then q is defined by

$$qj = \ln(1-L) \quad \Longrightarrow \quad q = \frac{\ln(1-L)}{\ln(1-x)} \tag{A.2}$$

Assume that at date t a disaster-news shock arrives. Let $\mathcal{I}_t$ be the information set just after arrival. Additionally, assume that expected returns follow an AR(1) process given by

$$\Delta\mathbb{E}_t\left[\ln R_{t+1+i}\right] = \phi^i\, \Delta\mathbb{E}_t\left[\ln R_{t+1}\right]$$

for $i \geq 0$ with $\phi \in [0,1)$. The news shock may affect dividend growth at time t, but is not persistent, that is

$$\Delta\mathbb{E}_t\left[\Delta d_{t+1+i}\right] = 0 \quad \text{for all } i \geq 1$$

As a result, iterating (A.1) forward, taking expectations conditional on $\mathcal{I}_t$, subtracting the pre-arrival counterpart, and summing the discounted differences yields the on-impact change in valuation

$$\Delta \ln \kappa_t \equiv \ln \kappa_t - \ln \kappa_t^- \approx \sum_{i\geq 0} \eta^i \left(\Delta\mathbb{E}_t\left[\Delta d_{t+1+i}\right] - \Delta\mathbb{E}_t\left[\ln R_{t+1+i}\right]\right) \tag{A.3}$$

where $\ln \kappa_t^-$ denotes the pre-arrival value. Inserting $\Delta\mathbb{E}_t\left[\ln R_{t+1+i}\right] = \phi^i\, \Delta\mathbb{E}_t\left[\ln R_{t+1}\right]$ and $\Delta\mathbb{E}_t\left[\Delta d_{t+1+i}\right] = 0$ for $i \geq 1$ into (A.3) yields

$$\Delta \ln \kappa_t \approx -\frac{\Delta\mathbb{E}_t\left[\ln R_{t+1}\right]}{1-\eta\phi} + \Delta\mathbb{E}_t\left[\Delta d_{t+1}\right] \tag{A.4}$$

If $\Delta\mathbb{E}_t\left[\Delta d_{t+1}\right] = 0$, i.e. no cash-flow news beyond the realized one-step drop j recorded at $t+1$, then

$$\Delta \ln \kappa_t \approx -\frac{\Delta\mathbb{E}_t\left[\ln R_{t+1}\right]}{1-\eta\phi}$$

Let the disaster at t be realized at $t+1$ with contraction $j < 0$. Furthermore, let the on-impact change in valuation be $\Delta \ln \kappa_t$. Then the co-jump in log returns at $t+1$ is given by

$$\Delta \ln R_{t+1} \equiv \ln R_{t+1} - \ln R_{t+1}^{-} \approx j + \Delta \ln \kappa_t$$

Applying the definition of q, given by (A.2), yields

$$qj = \ln(1-L) \quad \Longrightarrow \quad q = \frac{\Delta \ln R_{t+1}}{j} = 1 + \frac{\Delta \ln \kappa_t}{j}$$

Plugging back (A.4) gives

$$q \approx 1 - \frac{1}{j} \cdot \frac{\Delta \mathbb{E}_t \left[\ln R_{t+1}\right]}{1-\eta\phi} + \frac{\Delta \mathbb{E}_t \left[\Delta d_{t+1}\right]}{j} \tag{A.5}$$

If dividend-growth news is transient at impact, i.e. $\Delta \mathbb{E}_t \left[\Delta d_{t+1}\right] = 0$, then

$$q \approx 1 - \frac{1}{j} \cdot \frac{\Delta \mathbb{E}_t \left[\ln R_{t+1}\right]}{1-\eta\phi}$$

Simple inspection of (A.5) reveals that $q \geq 1$ if $\Delta \mathbb{E}_t \left[\ln R_{t+1}\right] \geq 0$ and $\Delta \mathbb{E}_t \left[\Delta d_{t+1}\right] \leq 0$. If either effect is strict, then $q > 1$ holds. Moreover, for a fixed $j < 0$, q is increasing in both the persistence ϕ and the size of the expected-return jump $\Delta \mathbb{E}_t \left[\ln R_{t+1}\right]$.

B Proof complements

B.1 Proof of Lemma 3

Proof. Define $J(\alpha) \equiv \ln\left(1-p+p\,e^{-\alpha j}\right)$. Then the log risk-free rate, given by (20), can be re-written as follows

$$\ln R_f = \delta + \gamma\mu - \frac{1}{2}\gamma^2\sigma^2 - \gamma\,\theta\,\sigma^2 + J(\theta) - J(\gamma+\theta)$$

Note that $\ln R_f$ is strictly decreasing in θ

$$\frac{\partial \ln R_f}{\partial \theta} = -\gamma\,\sigma^2 + J'(\theta) - J'(\gamma+\theta) < 0$$

as $J''(\cdot) > 0$ and thus $J'(\gamma+\theta) > J'(\theta)$. Furthermore, the cross-partial derivative is also negative

$$\frac{\partial^2 \ln R_f}{\partial\gamma\theta} = -\sigma^2 - J''(\gamma+\theta) < 0$$

since $J''(\cdot) > 0$. Subsequently, γ attains its highest admissible value if θ is zero. Define

$$f(\gamma) \equiv \delta + \gamma\mu - \frac{1}{2}\gamma^2\sigma^2 - J(\gamma) - \ln R_f = 0$$

Then f is strictly concave on $(0,\infty)$ as

$$f''(\gamma) = -\sigma^2 - J''(\gamma) < 0$$

with limits $\lim_{\gamma\downarrow 0} f(\gamma) = \delta - \ln R_f > 0$ and $\lim_{\gamma\to\infty} f(\gamma) = -\infty$. As a result, f has a unique root $\gamma^\dagger > 0$ via the Intermediate Value Theorem.

□

B.2 Proof of Proposition 4

Proof. Define

$$G(\gamma,\theta) \equiv \delta + \gamma\mu - \frac{1}{2}\gamma^2\sigma^2 - \gamma\,\theta\,\sigma^2 + J(\theta) - J(\gamma+\theta) - \ln R_f = 0$$

where $G(\gamma,\theta)$ again is strictly decreasing in θ

$$\frac{\partial G(\gamma,\theta)}{\partial\theta} = -\gamma\sigma^2 + J'(\theta) - J'(\gamma+\theta) < 0$$

as $J''(\cdot) > 0$ and thus $J'(\gamma+\theta) > J'(\theta)$. By Lemma 3 $G(\gamma,0) = f(\gamma) > 0$ for $\gamma \in (0,\,\gamma^\dagger)$ while

$$\lim_{\theta\to\infty} G(\gamma,\theta) = -\infty$$

as $J(\theta) - J(\gamma+\theta) \to \gamma j < 0$ for $\theta \to \infty$. Thus for each $\gamma \in (0, \gamma^{\dagger})$ there exists a unique $\theta(\gamma) > 0$ solving $G(\gamma,\theta) = 0$ via the Intermediate Value Theorem. The Implicit Function Theorem gives $\theta \in C^1$ with

$$\theta'(\gamma) = -\frac{\partial G(\gamma,\theta)}{\partial\gamma}\left[\frac{\partial G(\gamma,\theta)}{\partial\theta}\right]^{-1}$$

As already established, the denominator is strictly negative, while the numerator is given by

$$\frac{\partial G(\gamma,\theta)}{\partial\gamma} = \mu - \gamma\sigma^2 - \theta\sigma^2 - J'(\gamma+\theta)$$

Note that convexity of $J(\cdot)$ implies

$$J'(\gamma+\theta)\gamma \geq J(\gamma+\theta) - J(\theta)$$

Solving the weak inequality for $J'(\gamma+\theta)$ and substituting $J(\gamma+\theta) - J(\theta)$ from $G(\gamma,\theta)$ gives

$$J'(\gamma+\theta) \geq \frac{\delta - \ln R_f}{\gamma} + \mu - \frac{1}{2}\gamma\sigma^2 - \theta(\gamma)\,\sigma^2$$

Plugging this result back yields

$$\frac{\partial G(\gamma,\theta)}{\partial\gamma} \leq -\frac{1}{2}\gamma\sigma^2 - \frac{\delta - \ln R_f}{\gamma} < 0$$

as $\ln R_f < \delta$. Subsequently, the Implicit Function Theorem implies $\theta'(\gamma) < 0$. Last, note that as $\gamma \downarrow 0$ the mean-value identity $J(\gamma+\theta) - J(\theta) = J'(\zeta)\gamma$ with $\zeta \in (\theta, \theta+\gamma)$ and $J'(\zeta) \uparrow -j$ implies $\gamma\theta(\gamma)\,\sigma^2 \to \delta - \ln R_f$ and thus

$$\theta(\gamma) \sim \frac{\delta - \ln R_f}{\gamma\,\sigma^2} \quad \text{and} \quad \theta'(\gamma) \sim -\frac{\delta - \ln R_f}{\gamma^2\,\sigma^2}$$

□

B.3 Proof of Lemma 4

Proof. For $q \geq 1$ define $u \equiv e^{qj} \in (0,1)$ and $v \equiv e^{-(\gamma+\theta-q)j} = e^{-(\gamma+\theta)j}u > u$. Differentiating σ_r with respect to q gives

$$\sigma_r'(q) = \frac{S'(q)}{2\sqrt{S(q)}}$$

where $S(q)$ and $S'(q)$ are given by

$$S(q) = \ln\left(1+\frac{s^2}{m^2}\right) + 2\ln A(q) - \ln B(q)$$
$$S'(q) = 2pj\left[\frac{u}{1-p+pu} - \frac{u^2}{1-p+pu^2}\right]$$

Collecting the bracket in $S'(q)$ yields

$$\frac{u}{1-p+pu} - \frac{u^2}{1-p+pu^2} = \frac{(1-p)\,u(1-u)}{(1-p+pu)(1-p+pu^2)} > 0$$

for $u \in (0,1)$. Since $pj < 0$, it follows that $S'(q) < 0$. In the limit $S(q)$ is given by

$$\lim_{q\to\infty} S(q) = \ln\left(1+\frac{s^2}{m^2}\right) + \ln(1-p)$$

Solving for $\sigma_r(q) > 0$ as $q \to \infty$ yields the sufficient condition. If (31) is violated, the range of feasible q is restricted to $q \in [1, \bar{q})$, where $\bar{q}$ solves $S(\bar{q}) = 0$. This guarantees $\sigma_r(q) > 0$ on the feasible set.

□

B.4 Proof of Proposition 5

Proof. By the Implicit Function Theorem there is a unique $q(\gamma)$ solving $F(q,\gamma) = 0$ for each $\gamma \in \left(0,\ \gamma^{\dagger}\right)$. Totally differentiating the identity $F(q(\gamma), \gamma) \equiv 0$ yields

$$q'(\gamma) = -\left[\frac{\partial F(q,\gamma)}{\partial \gamma} + \theta'(\gamma)\frac{\partial F(q,\gamma)}{\partial \theta}\right]\left[\frac{\partial F(q,\gamma)}{\partial q}\right]^{-1}$$

Define $h(x) \equiv \frac{x}{1-p+px}$, which is increasing for $x > 0$. Differentiating (32) with respect to q gives

$$\frac{\partial F(q,\gamma)}{\partial q} = (\gamma+\theta)\,\rho\,\sigma\,\frac{S'(q)}{2\sqrt{S(q)}} + pj\,[h(u) - h(v)]$$

Lemma 4 shows that $S'(q) < 0$. Subsequently, the first term in $\frac{\partial F(q,\gamma)}{\partial q}$ is strictly negative. The second term satisfies $pj\,[h(u) - h(v)] > 0$ because h is increasing, $v > u$, and $j < 0$. Consequently, the sign of $\frac{\partial F(q,\gamma)}{\partial q}$ is ambiguous. Nevertheless, Assumption 1 and continuity of $\frac{\partial F(q,\gamma)}{\partial q}$ on the branch imply the sign is constant on each connected component. Differentiating (32) with respect to γ gives

$$\frac{\partial F(q,\gamma)}{\partial \gamma} + \theta'(\gamma)\frac{\partial F(q,\gamma)}{\partial \theta} = \left[1+\theta'(\gamma)\right]\left\{\rho\,\sigma\,\sigma_r(q) + J'(\gamma+\theta) - J'(\gamma+\theta-q)\right\}$$

where $J'(\gamma+\theta) - J'(\gamma+\theta-q) > 0$. Thus the curly bracket is strictly positive for $\rho > 0$. Since $\frac{\partial F(q,\gamma)}{\partial q}$ has constant sign it follows that

$$\text{sign}\left\{q'(\gamma)\right\} = -\varsigma \;\text{sign}\left\{1+\theta'(\gamma)\right\}$$

where $\varsigma \in \{+1,-1\}$ is the constant sign of $\frac{\partial F(q,\gamma)}{\partial q}$. By Proposition 4, $\theta'(\gamma) < 0$ and is continuous. Furthermore, if $\gamma \downarrow 0$, $\theta'(\gamma) \to -\infty$, implying $1+\theta'(\gamma) < 0$. Conversely, as $\gamma \uparrow \gamma^{\dagger}$

$$\lim_{\gamma\uparrow\gamma^{\dagger}} \left\{1+\theta'(\gamma)\right\} = \frac{-\mu - p\,j}{-\gamma\sigma^2 + J'(0) - J'(\gamma)} > 0$$

as the denominator is strictly negative and the numerator is strictly negative given $\mu > -p\,j$. Thus there exists a unique $\gamma^{\star} \in (0,\ \gamma^{\dagger})$ with $1+\theta'(\gamma^{\star}) = 0$ while $q(\gamma)$ is single-peaked on the admissible interval.

□

B.5 Proof of Proposition 6

Proof. Let $\omega(\gamma) = \gamma + \theta(\gamma)$. Since $j < 0$, $A(q) \to 1-p$ and $B(q) \to 1-p$ as $q \to \infty$. Subsequently, $\sigma_r(q) \to \sqrt{\ln\left(1+\frac{s^2}{m^2}\right) + \ln(1-p)}$ and $\ln\left[1-p+p\,e^{-(\omega(\gamma)-q)j}\right] \to \ln(1-p)$ as $q \to \infty$. As a result, the limit of (32) for $q \to \infty$ is given by

$$\begin{aligned}\lim_{q\to\infty} F(q,\gamma) \equiv\,& \delta + \gamma\mu - \frac{1}{2}\gamma^2\sigma^2 - \gamma\,\theta(\gamma)\,\sigma^2 \\ & + \omega(\gamma)\,\rho\,\sigma\sqrt{\ln\left(1+\frac{s^2}{m^2}\right) + \ln(1-p)} \\ & + \ln\left[1-p+p\,e^{-\theta(\gamma)\,j}\right] - \ln m\end{aligned} \tag{B.1}$$

For fixed γ, define $Q(\gamma) = \{q \geq 1 : F(q,\gamma) = 0\}$ and $\Gamma \equiv \{\gamma > 0 : Q(\gamma) \neq \emptyset\}$. Continuity of F in (q,γ) and the end-values $F(1,\gamma)$ and (B.1) ensure that $Q(\gamma) \neq \emptyset$ whenever those end-values have opposite signs (or one equals zero). Define $\gamma_{\min} := \inf\Gamma$. From Proposition 4, $\theta(\gamma) \to \infty$ as $\gamma \downarrow 0$, implying $\lim_{\gamma\downarrow 0}\lim_{q\to\infty} F(q,\gamma) = \infty$. Hence Γ is nonempty and $\gamma_{\min} > 0$.

Take any decreasing sequence $\gamma_n \downarrow \gamma_{\min}$ with $Q(\gamma_n) \neq \emptyset$, and pick $q_n \in Q(\gamma_n)$. The sequence $\{q_n\}$ is bounded: if $q_n \to \infty$ then $F(q_n,\gamma_n)$ converges to (B.1), which is greater than zero for n large, contradicting $F(q_n,\gamma_n) = 0$. Thus, along a subsequence, $q_n \to \tilde{q} \in [1,\infty)$ and by continuity $F(\tilde{q},\gamma_{\min}) = 0$. If $\tilde{q} > 1$, the point $(\tilde{q},\gamma_{\min})$ is interior to $\{q \geq 1\}$. By Assumption 1 and Proposition 5 the Implicit Function Theorem solves the local zero set $\{F = 0\}$ either as $q = q(\gamma)$ or as $\gamma = \gamma(q)$, producing solutions with $\gamma < \gamma_{\min}$, a contradiction. Hence $\tilde{q} = 1$ and

$$F\left(1, \gamma_{\min}\right)=0 \quad \text{and} \quad q\left(\gamma_{\min}\right)=1$$

Next, using (20) to eliminate the risk-free terms in (32) at $q=1$ gives the boundary identity

$$\begin{aligned} F(1, \gamma) = & \ln R_f - \ln m + \ln A(1) + (\gamma+\theta(\gamma))\, \rho\, \sigma\, \sigma_r(1) \\ & + J(\gamma+\theta(\gamma)) - J(\gamma+\theta(\gamma)-1) \end{aligned}$$

Differentiating yields

$$\frac{\partial F(1, \gamma)}{\partial \gamma} + \theta'(\gamma) \frac{\partial F(1, \gamma)}{\partial \theta} = \left[1+\theta'(\gamma)\right] \left\{\rho\, \sigma\, \sigma_r(1) + J'(\gamma+\theta) - J'(\gamma+\theta-1)\right\}$$

where as shown in Proposition 5 the curly bracket is strictly positive for $\rho > 0$. Proposition 5 gives $1+\theta'(\gamma) < 0$ near $\gamma \downarrow 0$, so $F(1, \gamma)$ is strictly decreasing there. Because $F\left(1, \gamma_{\min}\right) = 0$ and $F(1, \gamma) \rightarrow \infty$ as $\gamma \downarrow 0$ from $\theta(\gamma) \rightarrow \infty$, $\gamma_{\min}$ is the unique smallest root of $F(1, \gamma) = 0$. Note, that the same argument can be applied for $\gamma_{\max} := \sup \Gamma$ with the increasing sequence $\gamma_n \uparrow \gamma^{\dagger}$, which obtains γ^{+} as the unique largest root of $F(1, \gamma) = 0$.

Last, by Proposition 5 the risk-free equation (20) implies a unique risk-free upper bound $\gamma^{\dagger}$ at $\theta = 0$. Any feasible branch must satisfy $\gamma \leq \gamma^{\dagger}$, so the feasible set is an interval $\Gamma = \left(\gamma_{\min}, \gamma_{\max}\right)$ with $\gamma_{\max} \leq \gamma^{\dagger}$.

$\square$