FGV EPGE
ESCOLA BRASILEIRA DE ECONOMIA E FINANÇAS

Essays on Asset Pricing and Option Valuation

Gustavo Bulhões Carvalho da Paz Freire

Rio de Janeiro
2020

Gustavo Bulhões Carvalho da Paz Freire

Essays on Asset Pricing and Option Valuation

Tese submetida à Escola de Pós-Graduação em Economia como requisito parcial para a obtenção do grau de Doutor em Economia.

Fundação Getulio Vargas
Escola de Pós-Graduação em Economia

Orientador: Caio Almeida

Rio de Janeiro
2020

Dados Internacionais de Catalogação na Publicação (CIP)
Ficha catalográfica elaborada pelo Sistema de Bibliotecas/FGV

Freire, Gustavo Bulhões Carvalho da Paz
Essays on asset pricing and option valuation / Gustavo Bulhões Carvalho da Paz Freire. – 2020.
187 f.

Tese (doutorado) - Fundação Getulio Vargas, Escola Brasileira de Economia e Finanças.
Orientador: Caio Almeida.
Inclui bibliografia.

1. Modelo de precificação de ativos. 2. Mercado de opções. 3. Risco (Economia). I. Almeida, Caio Ibsen Rodrigues de. II. Fundação Getulio Vargas. Escola Brasileira de Economia e Finanças. III. Título.

CDD – 332.6

Elaborada por Márcia Nunes Bacha – CRB-7/4403

Gustavo Bulhões Carvalho da Paz Freire

"Essays on Asset Pricing and Option Valuation"

Tese apresentado(a) ao Curso de Doutorado em Economia do(a) EPGE Escola Brasileira de Economia e Finanças - FGV EPGE para obtenção do grau de Doutor(a) em Economia.

Data da defesa: 21/12/2020

ASSINATURA DOS MEMBROS DA BANCA EXAMINADORA

Presidente da Comissão Examinadora: Profº/ª Caio Ibsen Rodrigues de Almeida

Caio Ibsen Rodrigues de Almeida

Felipe Saraiva Iachan p/

Marcelo Fernandes p/

René Garcia p/

Ruy Ribeiro p/

Nos termos da Lei nº 13.979 de 06/02/20 - DOU nº 27 de 07/02/20 e Portaria MEC nº 544 de 16/06/20 - DOU nº 114 de 17/06/20 que dispõem sobre a suspensão temporária das atividades acadêmicas presenciais e a utilização de recursos tecnológicos face ao COVID-19, as apresentações das defesas de Tese e Dissertação, de forma excepcional, serão realizadas de forma remota e síncrona, incluindo-se nessa modalidade membros da banca e discente.

Lucas Jóver Maestri
Coordenador

Antonio de Araujo Freitas Junior
Pró-Reitor de Ensino, Pesquisa e Pós-Graduação FGV

Antonio Freitas, PhD
Pró-Reitor de Ensino, Pesquisa e Pós-Graduação
Fundação Getúlio Vargas

Instrução Normativa nº 01/19, de 09/07/19 - Pró-Reitoria FGV
Em caso de participação de Membro(s) da Banca Examinadora de forma não-presencial*, o Presidente da Comissão Examinadora assinará o documento como representante legal, delegado por esta I.N.
*Skype, Videoconferência, Apps de vídeo etc

Acknowledgments

First and foremost, I would like to thank my family, especially my parents, Fatima and Jorge, and my brothers, Gabriel and Guilherme, for all their support throughout my life. Special thanks also go to my girlfriend, Stefanie, whose encouragement, support and affection during the process of writing this thesis were of fundamental importance.

I am deeply grateful to my advisor, Professor Caio Almeida, for all his guidance and dedication, for countless enlightening discussions, and for showing me how research can be an exciting journey. I am also grateful to Professors Marcelo Fernandes, René Garcia, Felipe Iachan and Ruy Ribeiro for being on my thesis committee. I would further like to thank all professors at FGV EPGE for sharing their valuable knowledge, which were fundamental for my education. In particular, I thank Professors João Victor Issler and Marcelo Moreira, for which I had the pleasure to assist in their courses. I also thank Professor Markus Brunnermeier and the Department of Economics at Princeton University for their hospitality. Visiting the Bendheim Center for Finance added a lot to my research experience.

I would like to thank the many colleagues and friends at FGV EPGE and Princeton University, who helped make graduate school such an enriching experience. In particular, I am very grateful to João Thereze, for several stimulating discussions that contributed to this work. I have learned a lot from him and many other students over these years.

I also thank the amazing staff at FGV EPGE, for all the support they provide. In particular, I am grateful to Aline Neves for all the help during the process of visiting Princeton University.

Finally, I thankfully acknowledge financial support from: the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES), the Brazilian School of Economics and Finance at Fundação Getulio Vargas (FGV EPGE), and ANBIMA via the "XIV Prêmio ANBIMA de Mercado de Capitais".

"Perhaps I can best describe my experience of doing [research on] mathematics in terms of a journey through a dark unexplored mansion. You enter the first room of the mansion and it's completely dark. You stumble around bumping into the furniture, but gradually you learn where each piece of furniture is. Finally, after six months or so, you find the light switch, you turn it on, and suddenly it's all illuminated. You can see exactly where you were. Then you move into the next room and spend another six months in the dark. So each of these breakthroughs, while sometimes they're momentary, sometimes over a period of a day or two, they are the culmination of—and couldn't exist without—the many months of stumbling around in the dark that proceed them."

Andrew Wiles

Resumo

Esta tese é composta por quatro ensaios sobre precificação de ativos e opções. O primeiro ensaio estuda a precificação de opções de índice no contexto de mercados incompletos. Um novo método é fornecido para construir limites de preços para opções a partir de retornos do ativo subjacente e recuperar o investidor marginal que precifica uma determinada opção. Empiricamente, encontra-se que os preços das opções do S&P 500 são consistentes com os limites propostos. Opções com diferentes strikes no *cross-section* são precificadas por investidores heterogêneos que diferem essencialmente em sua avaliação do risco de eventos extremos negativos no mercado. O método fornece novos insights sobre a *expensiveness* relativa e os padrões de *skew* das opções de índice. O segundo ensaio deriva uma família de estimadores entrópicos generalizados para preços de opções usando apenas informações dos retornos do ativo subjacente. Uma aplicação empírica em grande escala identifica os estimadores que superam métodos de *benchmark* em termos de apreçamento fora da amostra. O terceiro ensaio propõe uma estrutura unificadora para a estimação do núcleo de apreçamento usando como informação os retornos do ativo subjacente e preços de opções. Esta estrutura permite recuperar o núcleo de apreçamento em função da riqueza endógena do investidor marginal nos mercados de índice e de opções de índice. Isso tem implicações importantes para o chamado *pricing kernel puzzle*. O quarto ensaio estima o risco de cauda para o Brasil e investiga as origens da variação do risco de cauda usando a cobertura de notícias da imprensa de negócios. Encontra-se que as preocupações com desastres são os principais impulsionadores da variação do risco de cauda e seus efeitos no mercado de ações e na economia real.

Palavras-chave: Apreçamento de Ativos; Apreçamento de Opções; Medida Neutra ao Risco; Núcleo de Apreçamento; Mercados Incompletos; Estimação Não-paramétrica; Risco de Cauda.

Abstract

This thesis consists of four essays on asset pricing and option valuation. The first essay studies the pricing of index options in the context of incomplete markets. A new method is provided to construct option price bounds from underlying returns and recover the marginal investor pricing a given option. Empirically, S&P 500 option prices are found to be consistent with the proposed bounds. Options with different strikes in the cross-section are priced by heterogeneous investors who essentially differ in their assessment of market downside risk. The method provides novel insights into the relative expensiveness and skew patterns of index options. The second essay derives a family of generalized entropic estimators for option prices using only information from underlying returns. A large-scale empirical application identifies the estimators outperforming benchmark methods in terms of out-of-sample pricing. The third essay proposes a unifying framework for the estimation of the pricing kernel using information from underlying returns and option prices. This framework allows to recover the pricing kernel as a function of the endogenous wealth of the marginal investor in the index and index option markets. This has important implications for the so-called pricing kernel puzzle. The fourth essay estimates tail risk for Brazil and investigates the origins of tail risk variation using news coverage from the business press. Disaster concerns are found to be the main drivers of tail risk variation and its effects on the stock market and the real economy.

Keywords: Asset Pricing; Option Pricing; Risk-Neutral Measure; Pricing Kernel; Incomplete Markets; Nonparametric Estimation; Tail Risk.

List of Figures

List of Tables

Contents

Introduction

This thesis consists of four essays addressing a number of important questions in the asset pricing and option valuation literature. The common thread in these essays stems from a methodology that was introduced by Almeida and Garcia (2017). They derive nonparametric moment restrictions to stochastic discount factors (SDFs), which are implied by SDFs minimizing a family of discrepancy functions. These restrictions generalize variance, entropy and higher-moment restrictions previously discovered in the literature, and can be used to diagnose asset pricing models taking into account higher-order moments of the candidate SDF. Of particular interest for this thesis is not the application to model diagnosis, but rather the fact that this method allows for the estimation of a family of nonnegative SDFs that correctly price a set of basis assets under no-arbitrage. The estimation procedure relies on a duality result, where the SDFs come from solutions of dual problems akin to maximizing concave utility functions. In addition, as these SDFs are nonnegative, each one produces a risk-neutral measure under which the returns on the basis assets are martingales.

The first essay extends the methodology above in terms of identifying and interpreting minimum dispersion risk-neutral measures estimated from returns on a given underlying asset. We particularly exploit the economic link that exists between the measures and investors maximizing a comprehensive class of well-behaved utility functions. In the context of incomplete markets, we characterize the set of admissible minimum dispersion measures and show how to construct option price bounds consistent with this set. Importantly, it is possible to recover the marginal risk-neutral measure (and the associated investor) pricing a given option. We use this method to reconcile option prices with underlying returns in incomplete markets. We find that S&P 500 option prices are mostly consistent with the proposed bounds, where options with different strikes in the cross-section are priced by heterogeneous investors who essentially differ in their assessment of market downside risk. Our method provides novel insights into the relative expensiveness and skew patterns of index options.

The second essay takes the perspective that each minimum dispersion risk-neutral measure estimated from underlying returns leads to a particular estimator to the price of an option. We suggest to consider this family of generalized entropic estimators to price options when the option market is not liquid enough to use standard methods that are option data-intensive. With Monte Carlo experiments, we show that, under sources of market incompleteness such as stochastic volatility and jumps, distinct estimators are necessary to price different options. In a large-scale empirical application, we compare the out-of-sample option pricing performance of the generalized entropic estimators with benchmark methods. We identify the estimators in the family that outperform the Black-Scholes and GARCH option models for different options in the cross-section.

The third essay provides a unifying framework for the literature that calculates empirical pricing kernels (EPKs) as the ratio of the option-implied state-price density and the historical return distribution of the market index. Such EPKs are often a U-shaped function of market returns, which is puzzling under the assumption of a complete market where the index proxies for wealth. We propose to estimate the pricing kernel by minimizing a convex discrepancy function subject to correctly pricing the underlying asset and observed options. The projection of our estimated pricing kernel onto market returns is the EPK identified by the literature. However, via the dual problem, we are also able to obtain the pricing kernel as a function of the endogenous wealth of the marginal investor pricing the index and index options. This implies that the U-shaped EPK can be rationalized as the projection of a pricing kernel that is monotonically decreasing in the endogenous wealth. That is, there is no puzzle when we recognize that options also constitute investment opportunities. The U-shaped pattern of the EPK arises because the marginal investor in the index and index options sells protection against large movements in the index.

The fourth essay follows Almeida et al. (2017) in using a minimum dispersion risk-neutral measure obtained from a set of portfolio returns to estimate tail risk. I estimate tail risk for Brazil and investigate the origins of tail risk variation. The tail risk measure peaks at stock market crashes, financial crises, political shocks and disaster events such as the coronavirus pandemic. Moreover, I find that tail risk has strong predictive power for market returns and negatively predicts real economic activity. In order to identify the investors' concerns associated to tail risk, I extract daily news from the largest financial newspaper in Brazil. The co-movement between news and tail risk indicates that tail risk variation and its effects on the stock market and the real economy are mainly driven by disaster concerns. These findings support recent models explaining asset pricing puzzles with time-varying disaster risk.

Chapter 1

Pricing of Index Options in Incomplete Markets

This paper provides a new method to reconcile option prices with underlying returns in incomplete markets. Using the returns on the underlying asset, we identify a set of risk-neutral measures associated to a comprehensive class of risk averse investors. From this set, we show how to construct option price bounds and recover the implied γ: a parameter that uniquely identifies the marginal investor pricing a given option. Empirically, we apply our method to nearly two million S&P 500 options. We find that option prices are mostly consistent with underlying returns, as the vast majority of prices lie within the bounds we calculate. Options with different strikes are priced by heterogeneous marginal investors who essentially differ in their assessment of market downside risk. This heterogeneity is time-varying and decreases during financial crises. Moreover, the implied γ allows us to interpret the structure of option prices taking into account information from the physical distribution. This provides novel insights into the relative expensiveness and skew patterns of index options before and after the 1987 crash.

This chapter is co-authored with Caio Almeida.

1.1 Introduction

Options are derivatives with payoffs dependent upon the stochastic evolution of the underlying asset price. A central question in empirical option pricing is whether investors rationally forecast the future asset price distribution when pricing options. In other words, whether the information coming from option prices in the cross-section is compatible with the one contained in the time series of underlying returns. Answering this question is challenging, as the risk-neutral distribution implied by option prices and the physical distribution of underlying asset values are not directly comparable without assumptions on investors' preferences.[1]

Assuming complete markets, a number of papers have investigated whether option prices are consistent with underlying returns. Ait-Sahalia and Lo (2000), Jackwerth (2000) and Rosenberg and Engle (2002) estimate the representative agent's preferences that reconcile the state-price density (SPD) extracted from option prices with the historical return distribution. The empirical pricing kernel obtained is non-decreasing, contradicting most economic models.[2] Without estimating preferences, Ait-Sahalia, Wang and Yared (2001) reject the null hypothesis that the SPD inferred from option prices and the SPD extracted from underlying returns under a one-factor structure are equal.

Arguably, a more plausible alternative in reconciling option prices and underlying returns is to consider incomplete markets.[3] The large trading volumes in option markets indeed suggest that options are non-redundant with respect to the underlying. In this context, the absence of arbitrage implies the existence of an infinity of pricing kernels or risk-neutral measures.[4] A way to deal with this multiplicity is to identify a set of pricing kernels from the underlying returns and ask if observed option prices are contained by the highest and lowest prices compatible with this set. Levy (1985) and Constantinides and Perrakis (2002) derive option price bounds consistent with monotonically decreasing pricing kernels. Cochrane and Saa-Requejo (2000) and Bernardo and Ledoit (2000) show how to construct bounds from stochastic discount factors (SDFs) with limited variance and gain-loss ratio, respectively. While price bounds are informative, they stop short in explaining *how* options are reconciled with the underlying return information.

In this paper, we provide a new method to reconcile option prices with underlying returns in incomplete markets. Using the returns on the underlying asset, we identify a set of risk-neutral measures associated to a comprehensive class of risk averse investors.

[1] See the discussion in Ait-Sahalia, Wang and Yared (2001).

[2] This is known as the pricing kernel puzzle. See Cuesdeanu and Jackwerth (2018) for a comprehensive survey on the puzzle.

[3] Ait-Sahalia and Lo (2000, p. 26) call for alternatives outside the assumption of complete markets.

[4] See Cochrane (2001). There is a direct correspondence between nonnegative pricing kernels and risk-neutral measures (see Section 1.2).

From this set, we show how to construct option price bounds and recover the implied γ: a parameter that uniquely identifies the marginal risk-neutral measure (and the associated investor) pricing a given option. This allows us to assess whether option prices are consistent with underlying returns and to recover the (potentially) heterogeneous marginal investors that price options in the cross-section. Moreover, similarly to the implied volatility, the implied γ can also be used to draw relative value comparisons between options, but accounting for the information coming from underlying returns.[5] We use our method to obtain novel empirical insights into the pricing of index options.

Our method starts from the premise that there is no *a priori* reason for a risk-neutral measure to deviate from the physical distribution other than correctly pricing the underlying returns. This suggests to consider risk-neutral measures that are the "closest" possible to the physical distribution by minimizing some notion of divergence. This premise has been adopted to obtain a unique minimum entropy measure.[6] Our novelty is to simultaneously select risk-neutral measures minimizing different discrepancies in a comprehensive family (Cressie and Read, 1984), extracting information from the whole set of measures.[7] Importantly, these measures are economically meaningful, in the sense that they can be obtained through dual portfolio problems for investors maximizing hyperbolic absolute risk aversion (HARA) utility functions.[8]

The risk-neutral measures minimizing Cressie-Read discrepancies are indexed by a single parameter $\gamma \in \mathbb{R}$. In particular, there is no solution to the minimum dispersion problem when $\gamma \to -\infty$ or $\gamma \to \infty$. We show that it suffices to search for an interval of γ's to explicitly identify the set of minimum dispersion measures. We characterize this set, which is endogenously determined by the underlying returns, and map the relation between the measures and options payoffs. Notably, the implied prices of any European option are monotonically decreasing in γ. Upper and lower option price bounds can then be calculated by the implied prices of the limiting measures of the interval of γ's. Furthermore, we can recover the implied γ, i.e., the parameter uniquely identifying the measure that correctly prices a given option.

Much like the implied volatility, the implied γ represents a monotonic mapping allowing for relative value comparisons between options in the cross-section. An implied volatility "smile" indicates deviations from the lognormal SPD associated to the Black-

[5]The implied volatility, a parameter that is inverted to characterize the Black and Scholes (1973) model consistent with a given option, is widely used as a mapping allowing for relative value comparisons between options with different strikes and maturities.

[6]See, for instance, Stutzer (1996) and Fujiwara and Miyahara (2003).

[7]A number of papers have considered SDFs minimizing the Cressie-Read family (Almeida and Garcia, 2017) or particular members of the family (Hansen and Jagannathan, 1991; Snow, 1991; Stutzer, 1995; Bansal and Lehmann, 1997; Liu, 2020) in order to diagnose asset pricing models.

[8]The class of HARA utility functions embeds as special cases several widely used utility functions, such as quadratic, exponential, logarithmic, power utility and constant relative risk aversion (CRRA).

Scholes model with at-the-money (ATM) implied volatility. In contrast, an implied γ smile would indicate deviations from the SPD that is closest to the underlying return physical distribution according to the implied γ of the ATM option. While the former denotes the misspecification of the Black-Scholes model, the latter is economically consistent with the existence of heterogeneous marginal investors with different attitudes towards risk. Investors with $\gamma < 1$ exhibit positive absolute prudence (Kimball, 1990), which is related to aversion to downside risk (Menezes, Geiss and Tressler, 1980) and to a convex marginal utility. Prudence is decreasing in γ, and for $\gamma > 1$ investors have negative prudence and concave marginal utilities.[9]

With simulated economies, we map the implications of different sources of market incompleteness to the relation between option prices and underlying returns.[10] The first economy is a Black-Scholes environment without intermediate trading, while the second one additionally incorporates stochastic volatility and jumps (Bates, 2000; Duffie, Pan and Singleton, 2000). While we find that option prices are consistent with underlying returns in both economies, there is a difference on how they are reconciled. In the Black-Scholes economy, where log-returns are Gaussian, the implied γ is flat across strikes and maturities, indicating the existence of a representative investor, in line with Rubinstein (1976).[11] Thus, the relative expensiveness of options in the cross-section is the same under the implied γ or the implied volatility. In the second economy, the additional sources of risk generate an implied γ smile, consistent with heterogeneous marginal investors with different risk preferences. In this environment, where log-returns are negatively skewed and leptokurtic, the implied γ takes into account information from the true physical distribution, while the implied volatility continues to denote deviations from Gaussian log-returns. This naturally leads to different conclusions about relative option expensiveness.

Empirically, we use our method to conduct a large-scale analysis of the pricing of S&P 500 options given the returns on the index. We select 1,998,832 options from January 2, 1986 to June 28, 2019. For each day, we estimate the conditional physical distribution from the underlying returns without assuming a parametric form. From this distribution, we calculate minimum dispersion price bounds for each option in our sample. For the option prices inside the bounds, we recover the implied γ.

We find that the vast majority of option prices lie within the minimum dispersion bounds (98.02% of the calls and 96.72% of the puts), indicating that option prices are

[9] As for risk aversion, all minimum dispersion risk-neutral measures are monotonically decreasing and consistent with positive absolute risk aversion.

[10] We consider economies coming from well-known parametric models. For each model, we calculate the minimum dispersion option price bounds using returns sampled from the model-implied physical distribution and compare them to option prices implied by risk-neutral parameters of the model.

[11] Rubinstein (1976) shows that the Black-Scholes option price is obtained without dynamic trading by specifying a CRRA utility function. Since CRRA is a particular case of HARA, the implied γ precisely identifies the representative investor associated with the stochastic process of underlying returns.

mostly consistent with underlying returns when we entertain the possibility that markets are incomplete. In particular, 97.53% of the short-term out-of-the-money (OTM) puts are contained by the bounds, contradicting the common notion that the left tail of the risk-neutral distribution is hard to be reconciled. On average, options are reconciled by an implied γ "smirk", where OTM puts (in-the-money (ITM) calls) are relatively more expensive than ITM puts (OTM calls). This indicates the existence of heterogeneous marginal investors in a segmented option market. OTM puts are priced by prudent investors averse to downside risk, consistent with the demand of such options as crash-insurance (Rubinstein, 1994). On the other hand, investors pricing OTM calls have negative prudence, in agreement with speculators buying them as a leveraged bet.

The heterogeneity of marginal risk-neutral measures is time-varying and provides novel insights into the relative expensiveness and skew patterns of index options before and after the 1987 crash. Although before the crash the implied volatility curve was flat, there was a reverse implied γ smirk indicating that OTM puts (and ITM calls) were actually cheaper, in relative terms, than ATM and ITM puts (and OTM calls). This is evidence that investors were pricing options using the Black-Scholes formula, even though this was inconsistent with the physical distribution of the underlying. In fact, if the physical distribution were lognormal, the implied γ curve would also have been flat.

After the crash, the structure of option prices changes to the usual smirk of the implied γ, where OTM puts (and ITM calls) become relatively more expensive in the cross-section. This confirms the "crashophobia" phenomenon (Rubinstein, 1994) when we consider information from the physical distribution, but to a lesser extent than suggested by the implied volatility smirk, as for several periods the implied γ's of low-strike and ATM options are close to each other. That is, under the implied γ, OTM puts are generally not too expensive. Furthermore, around financial crises, the implied γ gets approximately flat across strikes, even though there is an implied volatility smirk. In these periods, there is a single, prudent marginal investor in the cross-section.

Our paper is related to several strands of the literature. The first one is the already cited literature on option pricing in the context of incomplete markets, which mainly derives bounds on option prices. We contribute with a new approach that, besides generating meaningful option price bounds, allows for the recovery of the marginal risk-neutral measure (or SDF) pricing a given option. Also related is Gârleanu, Pedersen and Poteshman (2009), who model demand-pressure effects on option prices, where dealers trade options with end-users. While they are agnostic about the end-users' reasons for trade, our results are suggestive in linking the pricing of options with different strikes and maturities to marginal investors with different assessments of market downside risk.

Also previously cited is the literature addressing the issue of reconciling option prices with the time series of underlying returns. We provide evidence that option prices are rationally priced given the underlying returns when we consider incomplete markets.[12] Our results differ from those in Constantinides, Jackwerth and Perrakis (2009), who find substantial violations of stochastic dominance bounds in the S&P 500 option market. The main reason for the different results is that the conditional return distribution they estimate keeps the volatility constant over periods of up to three years, during which volatility varies considerably. In contrast, we adjust volatility daily in the estimation of the conditional return distribution. Moreover, we also contribute to this literature by identifying the risk-neutral measures and the associated marginal investors that reconcile each option with the underlying returns.

Our paper is also related to a growing literature using minimum dispersion risk-neutral measures (or SDFs) in empirical asset pricing (Stutzer, 1996; Ghosh, Julliard and Taylor, 2020; Korsaye, Quaini and Trojani, 2020; Kozak, Nagel and Santosh, 2020). We show how to identify the set of minimum dispersion measures, derive a number of interpretations for them and explore the information coming from the whole set of measures. In particular, while the literature usually focuses on the pricing kernels minimizing variance or entropy, our analysis highlights how SDFs minimizing distinct discrepancy functions can generate different implications to asset prices.

The remainder of the paper is organized as follows. Section 1.2 discusses the estimation of minimum dispersion risk-neutral measures and their properties. Section 1.3 describes how to identify the set of minimum dispersion measures and, from this set, construct option price bounds and recover the marginal measure pricing a given option. In Section 1.4, we explore with simulations the implications of different sources of risk and market incompleteness to the pricing of options. Section 1.5 discusses the findings of our empirical analysis on the relation between S&P 500 option prices and underlying returns. Section 1.6 concludes the paper.

1.2 Minimum Dispersion Risk-Neutral Measures

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and $\mathbf{R}$ denote a K-dimensional random vector on this space representing the gross returns of a set of K basis assets.[13] Given a risk-free rate R_f, denote by $\mathbf{R}^e \equiv \mathbf{R} - \mathbf{R}_f$ the vector of excess returns, where $\mathbf{R}_f \equiv R_f \mathbf{1}_K$ and $\mathbf{1}_K$ is a conformable vector of ones. In this setting, an admissible pricing kernel or stochastic discount factor (SDF) is a random variable m satisfying the pricing restrictions given by

[12] In the context of complete markets, Linn, Shive and Shumway (2018) and Barone-Adesi et al. (2020) show that the pricing kernel puzzle can be resolved by estimating the physical distribution in a forward-looking manner using option data.

[13] Our focus will be in the case that $K = 1$, but we begin with a more general setting.

the Euler equations:

$$\mathbb{E}(m\mathbf{R}^e) \equiv \int m\mathbf{R}^e \,\mathrm{d}\mathbb{P} = \mathbf{0}_K, \tag{1.1}$$

where $\mathbf{0}_K$ is a conformable vector of zeros and $\mathbb{P}$ is the physical probability measure.

There is a direct correspondence between nonnegative SDFs and risk-neutral measures. Dividing (1.1) by $\mathbb{E}(m)$ and considering the change of measure $\mathrm{d}\mathbb{Q} = \frac{m}{\mathbb{E}(m)}\,\mathrm{d}\mathbb{P}$, we get the following conditions:

$$\int \mathbf{R}^e \frac{m}{\mathbb{E}(m)}\,\mathrm{d}\mathbb{P} = \int \mathbf{R}^e \,\mathrm{d}\mathbb{Q} \equiv \mathbb{E}^{\mathbb{Q}}[\mathbf{R}^e] = \mathbf{0}_K, \tag{1.2}$$

where the risk-neutral measure $\mathbb{Q}$ is characterized by the state-price density $\mathrm{d}\mathbb{Q}$. A pricing kernel only produces a risk-neutral measure if it is nonnegative. Therefore, the set of admissible risk-neutral measures will be determined by the set of admissible nonnegative SDFs. In complete markets, the absence of arbitrage implies that a unique strictly positive admissible SDF exists. However, when the market is incomplete,[14] no-arbitrage implies an infinity of admissible strictly positive SDFs, and, hence, of admissible risk-neutral measures.

The information available for the estimation of a risk-neutral measure $\mathbb{Q}$ comes from the physical measure $\mathbb{P}$ and the moment conditions (1.2) defining the martingale property. In principle, there is no *a priori* reason for a risk-neutral measure to deviate from the physical distribution, other than satisfying (1.2). This suggests to consider risk-neutral measures that are the "closest" possible to the physical measure by minimizing some notion of divergence. Motivated by this premise, we consider the Cressie and Read (1984) family of discrepancy functions and, for each member of this family, choose the risk-neutral measure that is closest to $\mathbb{P}$ in the discrepancy sense.

The loss functions in the Cressie-Read family measuring the discrepancy between $\mathbb{Q}$ and $\mathbb{P}$ are indexed by a parameter $\gamma \in \mathbb{R}$ and defined by:

$$\phi_\gamma\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right) = \frac{\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)^{\gamma+1} - 1}{\gamma(\gamma+1)}. \tag{1.3}$$

Newey and Smith (2004) and Kitamura (2006) suggest using this family to measure divergence between distributions, while Almeida and Garcia (2017) adopt the Cressie-Read family to test asset pricing models taking into account higher moments of the admissible set of SDFs. Members of this family include the Euclidean divergence ($\gamma = 1$), the Kullback Leibler Information Criterion (KLIC) ($\gamma \to 0$), the Hellinger divergence ($\gamma = -1/2$), the empirical likelihood ($\gamma \to -1$) and Pearson's Chi-Square ($\gamma = -2$). As we show in

[14] In the one-period model setting considered here, the market is incomplete when the number of states of nature in Ω is larger than the number of basis assets K.

Appendix A 1.7.1, each discrepancy takes into account different sensitivities to higher moments of the state-price density. The Euclidean divergence is a quadratic loss function, measuring the variance between $\mathbb{Q}$ and $\mathbb{P}$. Discrepancies with γ close to one, as the KLIC, give small weights to higher moments as compared to the variance. The Pearson's Chi-Square gives equal (absolute) weight to all moments of the state-price density, while more negative γ's increasingly weight more, in absolute terms, higher moments.

A minimum dispersion risk-neutral measure solves the following minimization in the space of admissible measures with $I^{\phi_\gamma}(\mathbb{Q}, \mathbb{P}) \equiv \mathbb{E}\left[\phi_\gamma\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)\right] < \infty$:

$$\mathbb{Q}^* = \underset{Q}{arg\ min}\ \mathbb{E}\left[\phi_\gamma\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)\right] \equiv \int \phi_\gamma\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)\mathrm{d}\mathbb{P}, \ \text{ s.t. } \ \mathbb{E}^{\mathbb{Q}}\left[\mathbf{R}^e\right] = \mathbf{0}_K, \tag{1.4}$$

where, under the restriction that $\mathbb{Q}$ is absolutely continuous with respect to $\mathbb{P}$, $\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}$ is a Radon-Nikodym derivative[15]. The risk-neutral measure must also be nonnegative and integrate to one. By definition, we have that any admissible measure $\mathbb{Q}$ satisfies $I^{\phi_\gamma}(\mathbb{Q}, \mathbb{P}) \geq I^{\phi_\gamma}(\mathbb{Q}^*, \mathbb{P})$. Moreover, $I^{\phi_\gamma}(\mathbb{Q}, \mathbb{P}) \geq 0$, with equality only when $\mathbb{Q} = \mathbb{P}$. Therefore, we have $I^{\phi_\gamma}(\mathbb{Q}^*, \mathbb{P}) = 0$ only in the case that $\mathbb{P}$ satisfies the risk neutrality constraint in (1.4), that is, when all mean excess returns are zero or, equivalently, that there is no risk premia in the market defined by the basis assets.

The minimum dispersion problem (1.4) treats the physical distribution nonparametrically by comparing the empirical distribution of the returns with the family of distributions implied by all probability measures satisfying the conditions of risk neutrality. This problem is of infinite dimension and difficult to solve. In the Online Appendix 1.12, we follow Kitamura (2006) and Almeida and Garcia (2017) and make use of results in Borwein and Lewis (1991) to show (in Theorem 1) that it can be solved via a much simpler finite dimensional dual problem. Absence of arbitrage is a fundamental condition for the solutions of the primal and dual problems to coincide.

In order to estimate minimum dispersion risk-neutral measures from data on returns, we consider the sample version of problem (1.4). In this case, the sample space Ω is discrete and finite, with states of nature $k = \{1, ..., n\}$, where $n > K$. Let $\{\mathbf{R}_k\}_{k=1}^n$ be the observed gross returns of the K basis assets, where each $\mathbf{R}_k$ is independent and identically distributed according to $\mathbb{P}$. The unknown physical measure $\mathbb{P}$ can be replaced by the empirical measure $\mathbb{P}_n$ that gives weights $\pi_k = 1/n$ to the realization of each state of nature.[16] This allows us to exchange the expectation $\mathbb{E}$ with its sample counterpart $\frac{1}{n}\sum_{k=1}^n \equiv \sum_{k=1}^n \pi_k$. In the following corollary implied by Theorem 1 in the Online Ap-

---

[15] i.e., $\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} : \Omega \to [0, \infty)$ is a measurable function such that for any measurable set $A \subseteq \Omega$, $\mathbb{Q}(A) = \int_A \frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\mathrm{d}\mathbb{P}$.

[16] This constitutes an optimal nonparametric estimator for $\mathbb{P}$. For more details, see Kitamura (2006).

pendix 1.12, we summarize the results for the sample version of the problem of finding a minimum dispersion Cressie-Read risk-neutral measure:

Corollary 1. *Consider the primal problem:*

$$\begin{aligned} &\underset{\{\pi_1^Q,\dots,\pi_n^Q\}}{min} \quad \sum_{k=1}^n \pi_k \frac{\left(\pi_k^{\mathbb{Q}}/\pi_k\right)^{\gamma+1}-1}{\gamma(\gamma+1)}, \\ &s.t. \;\; \sum_{k=1}^n \pi_k^{\mathbb{Q}} \left(\boldsymbol{R}_k - \boldsymbol{R}_f\right) = \boldsymbol{0}_K, \;\; \sum_{k=1}^n \pi_k^{\mathbb{Q}} = 1, \; \pi_k^{\mathbb{Q}} \geq 0 \;\; \forall\, k. \end{aligned} \tag{1.5}$$

Absence of arbitrage in the observed sample implies that the value of the primal problem coincides (with dual attainment) with the value of the dual problem below:

i) if $\gamma > 0$:

$$\lambda_\gamma^* = \underset{\lambda \in R^K}{arg\ max} - \frac{1}{\gamma+1} \sum_{k=1}^n \pi_k \left(1 + \gamma\lambda' \left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)\right)^{\left(\frac{\gamma+1}{\gamma}\right)} I_{\Lambda_\gamma(\boldsymbol{R}_k)}(\lambda), \tag{1.6}$$

ii) if $\gamma < 0$:

$$\lambda_\gamma^* = \underset{\lambda \in \Lambda_\gamma}{arg\ max} - \frac{1}{\gamma+1} \sum_{k=1}^n \pi_k \left(1 + \gamma\lambda' \left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)\right)^{\left(\frac{\gamma+1}{\gamma}\right)}, \tag{1.7}$$

iii) if $\gamma = 0$, the maximization is unconstrained:

$$\lambda_0^* = \underset{\lambda \in R^K}{arg\ max} \; - \sum_{k=1}^n \pi_k \, e^{\lambda' \left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)}, \tag{1.8}$$

where $\Lambda_\gamma = \{\lambda \in \mathbb{R}^K : for\ k = 1, ..., n,\ (1 + \gamma\lambda' (\boldsymbol{R}_k - \boldsymbol{R}_f)) > 0\}$, $\Lambda_\gamma(\boldsymbol{R}_k) = \{\lambda \in \mathbb{R}^K : (1 + \gamma\lambda' (\boldsymbol{R}_k - \boldsymbol{R}_f)) > 0\}$ and $I_A(x) = 1$ if $x \in A$, and 0 otherwise.

The minimum dispersion risk-neutral measure can then be recovered from the first-order conditions of (1.6), (1.7) and (1.8) with respect to λ, evaluated at λ_γ^*:

$$\pi_k^{\mathbb{Q}*}(\gamma, \boldsymbol{R}) = \frac{(1 + \gamma\lambda_\gamma^{*\prime}(\boldsymbol{R}_k - \boldsymbol{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\boldsymbol{R}_k)}(\lambda_\gamma^*)}{\sum_{i=1}^n (1 + \gamma\lambda_\gamma^{*\prime}(\boldsymbol{R}_i - \boldsymbol{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\boldsymbol{R}_i)}(\lambda_\gamma^*)}, \; k = 1, ..., n;\ \gamma > 0, \tag{1.9}$$

$$\pi_k^{\mathbb{Q}*}(\gamma, \boldsymbol{R}) = \frac{\left(1 + \gamma\lambda_\gamma^{*\prime} \left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)\right)^{\frac{1}{\gamma}}}{\sum_{i=1}^n \left(1 + \gamma\lambda_\gamma^{*\prime} \left(\boldsymbol{R}_i - \boldsymbol{R}_f\right)\right)^{\frac{1}{\gamma}}}, \; k = 1, ..., n;\ \gamma < 0, \tag{1.10}$$

$$\pi_k^{\mathbb{Q}*}(0, \boldsymbol{R}) = \frac{e^{\lambda_0^*\left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)}}{\sum_{i=1}^n e^{\lambda_0^*\left(\boldsymbol{R}_i - \boldsymbol{R}_f\right)}}, \; k = 1, ..., n;\ \gamma = 0. \tag{1.11}$$

Proof. See Online Appendix 1.12. □

For each discrepancy γ, the solution λ_γ^* of the dual problem leads to a different minimum dispersion risk-neutral measure. For $\gamma = 0$, $\pi^{\mathbb{Q}*}$ is exponential, hence strictly positive. For $\gamma < 0$, we search for the Lagrange multiplier λ in Λ_γ, implying that the risk-neutral measure is strictly positive when the constrained maximization has a solution. For $\gamma > 0$, the implied measure will be strictly positive, unless there are states of nature k such that $(1 + \gamma\lambda_\gamma^*(\mathbf{R}_k - \mathbf{R}_f)) \leq 0$. In this case, the indicator function $I_{\Lambda_\gamma(\mathbf{R}_k)}$ will set to zero the weights corresponding to these states.

1.2.1 Economic Interpretation

In addition to their statistical motivation, the minimum dispersion risk-neutral measures are also economically meaningful. The dual problem coming from each Cressie-Read discrepancy minimization can be economically interpreted as a standard optimal portfolio problem for an investor with HARA utility:

$$u^\gamma(W) = -\frac{1}{\gamma+1}(b - a\gamma W)^{\frac{\gamma+1}{\gamma}}, \tag{1.12}$$

with $a > 0$ and $b - a\gamma W > 0$, which guarantees that the function u^γ is well-defined, concave and strictly increasing. More specifically, the investor distributes her initial wealth W_0 putting $\tilde{\lambda}_j$ units of wealth on the risky asset R_j and the remaining $W_0 - \sum_{j=1}^K \tilde{\lambda}_j$ in a risk-free asset paying R_f. Terminal wealth is then given by $W(\tilde{\lambda}) = W_0 R_f + \sum_{j=1}^K \tilde{\lambda}_j (R_j - R_f)$ and the investor maximizes expected utility:

$$\tilde{\lambda}_\gamma^* = \underset{\tilde{\lambda} \in R^K}{max}\ \mathbb{E}\left[u^\gamma(W(\tilde{\lambda}))\right]. \tag{1.13}$$

In the Online Appendix 1.12, we show that solving (1.13) is equivalent to solving the population version of the dual problems (1.6), (1.7) and (1.8). In particular, letting λ_γ^* denote the dual problem solution, $\tilde{\lambda}_\gamma^* = -\lambda_\gamma^*(b - a\gamma W_0 R_f)/a$ if $\gamma \neq 0$ and $\tilde{\lambda}_\gamma^* = -\lambda_\gamma^*/a$ if $\gamma = 0$. By solving the dual problem, we obtain the optimal portfolio solution for all values of W_0, where $\tilde{\lambda}_\gamma^*$ is proportional to λ_γ^* and has opposite sign.

The result above essentially shows that the minimum dispersion risk-neutral measures can be interpreted as the marginal utilities of HARA investors. The HARA class constitutes a considerably general class of risk averse investors, encompassing as special cases several widely used utility functions. If $\gamma = 1$, it specializes to quadratic utility, while $\gamma \to 0$ yields exponential utility. The power utility function is obtained if $\gamma < 0$ and $a = -1/\gamma$, specializing to CRRA utility if, further, $b = 0$. Logarithmic utility obtains for $a = 1$ and $\gamma \to -1$, while linear utility (the risk-neutral case) is achieved when $\gamma \to -\infty$.

The parameter γ also indexes important notions of attitude towards risk, such as the Arrow-Pratt coefficient of absolute risk aversion and the absolute prudence of Kimball (1990). These are given by $A^\gamma(W) = -u^{\gamma\prime\prime}(W)/u^{\gamma\prime}(W) = a/(b - \gamma a W)$ and

$P^\gamma(W) = -u^{\gamma\prime\prime\prime}(W)/u^{\gamma\prime\prime}(W) = a(1-\gamma)/(b-\gamma aW)$, respectively. To help visualize how the risk preferences change across HARA investors, Figure 1.1 plots $A^\gamma(W)$ and $P^\gamma(W)$ for different values of W and γ. Absolute risk aversion is always positive, characterizing risk averse behavior.[17] The main difference between the investors is regarding the absolute prudence, which is positive for $\gamma < 1$, zero for $\gamma = 1$ and negative for $\gamma > 1$. Positive prudence is related to aversion to downside risk (Menezes, Geiss and Tressler, 1980) and to a convex marginal utility, while negative prudence to a concave marginal utility. Moreover, as can be seen, the smaller the γ the more prudent is the investor.

1.2.2 Comparing Minimum Dispersion Risk-Neutral Measures

Each minimum dispersion risk-neutral measure is the smallest modification to $\mathbb{P}$ that correctly prices the basis assets, according to a different discrepancy function. That is, each measure distorts the physical measure in a particular way. This is illustrated in the left panel of Figure 1.2 for different γ's, with measures obtained from lognormal returns on a single risky basis asset, coming from the Black-Scholes model (see Appendix B 1.8). The risk-neutral measures give more weight to negative excess returns and less to positive ones compared to the empirical probabilities, which weight equally the returns. All measures are monotonically decreasing in the returns, but their curvatures depend on how each discrepancy weights higher-order moments of the state-price density. For $\gamma = 1$, the Cressie-Read discrepancy is quadratic, implying that the risk-neutral measure is linear. For values of γ smaller than 1, the measures become more convex, increasingly weighting more extreme returns. In contrast, for $\gamma > 1$ the measures are concave, increasing the weights in moderate returns. Moreover, as γ increases, the measures start to give zero weight for the most positive returns, due to the indicator function $I_{\Lambda_\gamma(R_k)}$.

The shapes of the risk-neutral measures can also be explained by the economic interpretation of the dual problem. The weight π_k that a particular measure gives to state k represents the willingness to pay for the state claim paying one in this state and zero elsewhere. The physical measure, for instance, is associated to a risk-neutral investor with linear utility ($\gamma \to -\infty$), implying constant marginal utility. Except for this case, HARA investors are risk averse, with decreasing marginal utility: a state claim that pays off when wealth is low has a relatively high price, because it allows the investor to smooth wealth across future states of nature. On the other hand, the absolute prudence of the investors determines the curvature of their marginal utilities. For $\gamma = 1$, the prudence is zero and the risk-neutral measure is linear. For $\gamma > 1$, the absolute prudence is negative, implying concave marginal utilities, while $\gamma < 1$ is associated to positive absolute prudence

[17]For a fixed W, investors with larger γ's have higher risk aversion. This difference is exacerbated at higher levels of wealth, since $\gamma > 0$ implies increasing absolute risk aversion (i.e., risk aversion increases if wealth increases), while $\gamma < 0$ characterizes decreasing absolute risk aversion.

and convex risk-neutral measures. The smaller the γ, the more prudent the investor is, willing to pay more for state claims that pay off in the "worst" states of nature, due to the aversion to downside risk.[18]

Yet another way of depicting a minimum dispersion risk-neutral measure is through the corresponding state-price density (SPD). The SPD is a density function assigning risk-neutral probabilities to the future possible values of the asset price. For a given γ, the SPD can be approximated by the histogram of returns (times the current price of the asset) under the minimum dispersion risk-neutral probabilities. Similarly, the physical distribution is approximated by the histogram of returns under the empirical measure $1/n$. This is illustrated in the right panel of Figure 1.2, where we compare the SPD implied by $\gamma = -2$ with the histogram of future prices coming from the physical distribution, assuming a current price of 100. Under the minimum dispersion SPD, smaller future prices for the asset are more likely to occur than under the physical distribution, while the opposite is true for higher future prices. That is, the minimum dispersion SPD is more skewed to the left.

1.3 Price Bounds and Marginal Risk-Neutral Measures

In order to provide meaningful option price bounds in incomplete markets and be able to recover the marginal risk-neutral measures in the cross-section of options, we propose to restrict the admissible set to the minimum dispersion risk-neutral measures. To derive our method, the first natural question is if the minimum dispersion measures should cover the whole Cressie-Read family. To answer this, we investigate what happens to the minimum discrepancy problem when $\gamma \to -\infty$ or $\gamma \to \infty$. Below we show that there is no solution to the problem in these limiting cases.

Proposition 1. *There is no solution to the minimum dispersion problem (1.5) when $\gamma \to -\infty$ or $\gamma \to \infty$.*

Proof. See Online Appendix 1.12. □

Therefore, it is necessary to identify a limited set of γ's to define the set of minimum dispersion risk-neutral measures. We focus on characterizing this set from a single asset

[18] The shapes of the risk-neutral measures are robust to the returns considered. That is, regardless of the distribution generating the returns, the minimum dispersion risk-neutral measures will be monotonically decreasing, convex for $\gamma < 1$, linear for $\gamma = 1$ and concave for $\gamma > 1$. Moreover, the convexity will always decrease with γ. What may differ is that, for a fixed value of γ, the implied risk-neutral measure can be more or less convex depending on the return distribution being more or less skewed to the left.

(the underlying). In the Online Appendix 1.12, we show how our method can be applied to the general case of deriving price bounds for a non-redundant asset with respect to K basis assets.

We proceed by characterizing when the dual problem has a solution. While any solution will be nonnegative, we are also interested in identifying when the solution will be strictly positive.[19] For $\gamma = 0$, the dual problem is unconstrained and the implied measure is exponential, therefore strictly positive. Negative γ's imply a constrained maximization problem, which may not have a solution. On the other hand, positive γ's generate measures with indicator functions that can set weights to zero in some states of nature. In the next proposition, we show that the minimum dispersion risk-neutral measures are monotonic on the properties above. That is, as γ becomes more negative, the constraint is satisfied until a given minimum negative γ is reached, while as γ becomes larger, the measures will be strictly positive until a given positive γ is reached.

Proposition 2. *Let $\mathbb{E}(R - R_f) > 0$ and λ^*_γ be the solution for a given γ to the dual problem in Corollary 1 when there is just one risky basis asset.*[20] *Let also max_i denote the i^{th} highest value.*

(i) Let $\gamma < 0$ in a neighborhood of zero. Then $\gamma\lambda^*_\gamma > 0$ and, as γ becomes more negative, $\gamma\lambda^*_\gamma$ increases and the dual problem only has a solution while $\gamma\lambda^*_\gamma < -1/min\{R_k - R_f\}$.

(ii) Let $\gamma > 0$ in a neighborhood of zero. Then $\gamma\lambda^*_\gamma < 0$ and, as γ increases, $\gamma\lambda^*_\gamma$ decreases and the implied risk-neutral measure is strictly positive while $\gamma\lambda^*_\gamma > -1/max\{R_k - R_f\}$. When $\gamma\lambda^*_\gamma \leq -1/max\{R_k - R_f\}$, the measure is not strictly positive anymore. The i^{th} zero in the measure will appear when $\gamma\lambda^*_\gamma \leq -1/max_i\{R_k - R_f\}$.

Proof. See Online Appendix 1.12. □

The proposition above provides an algorithm to find the admissible set of minimum dispersion risk-neutral measures from the underlying returns $\{R_k\}_{k=1}^n$. Starting from $\gamma = 0$, one can solve the dual problem for a grid of decreasing γ's becoming more negative, keeping track of the value of $\gamma\lambda^*_\gamma$. The $\underline{\gamma}$ where $\gamma\lambda^*_\gamma$ gets arbitrarily close to $-1/\min\{R_k - R_f\}$ will be the limiting measure of the set for the left hand side of the family. Analogously, one can solve the dual problem for a grid of increasing positive γ's, until $\gamma\lambda^*_\gamma$ gets smaller than $-1/\max\{R_k - R_f\}$. The $\overline{\gamma}$ where this happens will be the

---

[19] While any measure solving the minimum dispersion dual problem will be consistent with no-arbitrage in the market spanned by the basis assets, measures that are not strictly positive may generate arbitrage opportunities when we consider pricing non-redundant assets.

[20] $\mathbb{E}(R - R_f) > 0$ represents the usual case where the risk premium of the underlying asset is positive. See the Online Appendix 1.12 for the case where $\mathbb{E}(R - R_f) < 0$. The results are analogous.

limiting measure for the right hand side of the family.[21] The interval $[\underline{\gamma}, \overline{\gamma}]$ characterizes the set of strictly positive minimum dispersion risk-neutral measures and is endogenously determined by the underlying returns. It is also possible to include measures with zeros in some states of nature by allowing for γ's greater than $\overline{\gamma}$. In fact, it is possible to choose the exact number of zeros allowed for the risk-neutral measure. For now, we shall focus on the strictly positive set $[\underline{\gamma}, \overline{\gamma}]$.[22]

Since $[\underline{\gamma}, \overline{\gamma}]$ explicitly identifies all the measures within the set, given a grid it is possible to calculate the implied price of a non-redundant asset by each measure. Moreover, because the risk-neutral measures are continuous functions of γ, the implied prices will change continuously with γ and reach a maximum and minimum in the interval. One can then calculate upper and lower price bounds for a non-redundant asset by taking the maximum and minimum prices implied by the measures in $[\underline{\gamma}, \overline{\gamma}]$. More specifically, let x be the payoff we want to price. The lower minimum dispersion price bound solves:

$$\underline{C} = \min_{\{\pi_\gamma^{\mathbb{Q}}\}} \frac{1}{R_f} \sum_{k=1}^{n} \pi_{\gamma k}^{\mathbb{Q}} x_k, \quad s.t. \;\; \gamma \in [\underline{\gamma}, \overline{\gamma}], \tag{1.14}$$

where the upper bound $\overline{C}$ solves the corresponding maximum. If, *a posteriori*, the asset price C is observed and $C \in [\underline{C}, \overline{C}]$, one can recover the risk-neutral measure (and the implied γ) that correctly prices the asset.

We focus on the particular case where x is the payoff of a European option. Considering a call option expiring in T days, its price is given by:

$$C = \mathbb{E}^{\mathbb{Q}} \left[\frac{\max\left(p_0 R - X, 0\right)}{(1 + r_f)^T} \right], \tag{1.15}$$

where $\mathbb{E}^{\mathbb{Q}}$ is the expectation taken with respect to a risk-neutral measure $\mathbb{Q}$, X is the option strike, r_f is the daily risk-free rate, p_0 is the current price of the underlying asset, R is the T-day return from the physical distribution and $p_0 R \equiv p_T$ is the price of the underlying at time T. The sample counterpart to (1.15) is given by an expectation over $k = 1, ..., n$ states of nature, for which we have returns R_k drawn from the physical distribution with weights $\pi_k = 1/n$:

[21]For $\gamma < 0$, the problem is constrained by Λ_γ. Computationally, this constraint is imposed and will not be violated. In practice, what happens is that the dual problem stops having a solution when $\gamma\lambda_\gamma^*$ gets arbitrarily close to $-1/\min\{R_k - R_f\}$, up to the fifth decimal place. Solutions beyond this point will be associated to non-negligible pricing errors for the underlying asset, which is the computational counterpart of the dual problem not having a solution. For $\gamma > 0$, $\gamma\lambda_\gamma^* > -1/\max\{R_k - R_f\}$ is not a constraint to the problem, so it will happen that $\gamma\lambda_\gamma^* \leq -1/\max\{R_k - R_f\}$ for some γ in the increasing sequence of γ's. An appropriate grid of γ's is with 0.1 spacing.

[22]The remaining analysis in this section is completely analogous for the case allowing for risk-neutral measures with zeros, which yields a larger interval $[\underline{\gamma}, \overline{\gamma}']$, where $\overline{\gamma}' > \overline{\gamma}$.

$$C = \sum_{k=1}^{n} \pi_k^{\mathbb{Q}} \left[\frac{\max\left(p_0 R_k - X, 0\right)}{(1 + r_f)^T} \right], \tag{1.16}$$

where $\pi_k^{\mathbb{Q}}$ is the risk-neutral counterpart of the empirical measure π_k and must correctly price the underlying returns (i.e., satisfy the martingale property). As the number of states of nature n is larger than the number of basis assets, the market is incomplete and there exists multiple risk-neutral measures. In this context, the underlying returns R_k and the risk-free rate given by $R_f = (1 + r_f)^T$ can be used to obtain the interval $[\underline{\gamma}, \overline{\gamma}]$ of minimum dispersion risk-neutral measures and calculate price bounds according to (1.14) with $x_k = \max\left(p_0 R_k - X, 0\right)$.

Given the convexity of option payoffs, the option prices implied by the minimum dispersion measures will be monotonically decreasing in γ. This is because extreme positive (negative) returns are the ones that make the call (put) options pay at expiration, and, as the left panel of Figure 1.2 shows, the convexity of the measures (or the weights given to the extreme returns) decreases as γ increases. This implies that the calculation of option price bounds from the minimum dispersion admissible set will be even simpler: the limiting measures of the interval $[\underline{\gamma}, \overline{\gamma}]$ will be the ones defining the upper and lower price bounds, respectively.

This is illustrated in Figure 1.3 for an ATM call option with one month to maturity. The minimum dispersion admissible set is calculated using 1-month underlying returns from the SVCJ model (see Appendix B 1.8) and is depicted by the interval $[\underline{\gamma}, \overline{\gamma}]$. The curve in blue plots the option prices implied by each risk-neutral measure in $[\underline{\gamma}, \overline{\gamma}]$. Since the implied prices decrease monotonically in γ, the limiting measures of the admissible set determine the price bounds. Moreover, as the true option price implied by the SVCJ model (in red) is contained by the bounds, it is possible to compute the associated unique implied γ. In this case, the implied γ is equal to -0.8, indicating that the marginal risk-neutral measure is convex and associated to a prudent investor.

The monotonicity of implied option prices with respect to γ also allows the implied γ to be seen as a measure of relative option expensiveness that nonparametrically takes into account information from the underlying return physical distribution. Considering a cross-section of options, a smile of the implied γ would indicate deviations from the minimum dispersion risk-neutral measure that correctly prices the ATM option. Options with smaller (higher) implied γ than the ATM implied γ require larger (smaller) weights in extreme underlying returns from the physical distribution, being relatively more expensive (cheaper). While the implied γ provides relative value comparisons between options using the information from the physical distribution, the implied volatility draws comparisons assuming that underlying returns are lognormal.

In Section 1.4, we explore in detail the implications of our approach to option prices in simulated economies accounting for different sources of market incompleteness. In Section 1.5, we use our method to obtain new empirical results on the relation between the pricing of index options and underlying returns. In the next subsection, we discuss how our approach relates to other methods that impose additional restrictions beyond no-arbitrage to reduce the set of admissible measures and construct price bounds. One important difference from these methods is that we explicitly identify all the risk-neutral measures inside the restricted admissible set. This allows us to map the measures and risk preferences of marginal investors that are consistent with observed prices of options.

1.3.1 Comparison with Other Methods

Good-Deal Pricing - Cochrane and Saa-Requejo (2000)

In order to restrict the set of admissible SDFs and learn about the value of a non-redundant asset with payoff x, Cochrane and Saa-Requejo (2000) require the pricing kernel to be nonnegative and have volatility lower than a given exogenous upper limit A. By duality, this is equivalent to an upper limit to the (arbitrage-adjusted) Sharpe ratio, restricting the existence of good-deals. The smallest feasible A will be given by the nonnegative SDF with minimum variance x^*. Therefore, as one exogenously varies A, the bounds to x can range from the exact price implied by x^* to the no-arbitrage bounds. Note that, for a given A, not all SDFs in the no-good-deal admissible set will be economically well-behaved.[23] Furthermore, it is necessary to choose A *a priori*.

In this framework, our approach can be viewed as a "well-behaved" good-deal pricing. To see this, note that x^* is the SDF obtained by the Cressie-Read minimization with $\gamma = 1$.[24] If x^* is inside the minimum dispersion admissible set, all other risk-neutral measures in our set have higher variance, so they naturally impose an endogenous maximum limit $\tilde{A}$ to the SDF volatility and, via duality, to the Sharpe ratio. The difference is that if we used $\tilde{A}$ in the method of Cochrane and Saa-Requejo, the no-good-deal admissible set would be much larger, while we identify a smaller subset of measures that correspond to marginal utilities of HARA investors.

In addition, by a similar argument to that of variance, since each element in our admissible set is a minimum dispersion measure, our set endogenously bounds the SDF discrepancy according to each γ inside the set. By duality, as we show in Appendix A 1.7.2, this is equivalent to limiting the maximum generalized Sharpe ratio attainable in the extended market with options. From the generalized no-good-deal admissible set, we select only measures consistent with HARA investors.

[23] While the covariance between the basis returns and the SDF must be negative in their method, this does not imply global monotonicity of the pricing kernel.

[24] As they impose that $m \geq 0$, we can refer interchangeably to the SDF and the risk-neutral measure.

Good-Deal Pricing - Bernando and Ledoit (2000)

Bernardo and Ledoit (2000) start from a benchmark investor implying an SDF m^*, which in turn defines the gain-loss ratio as a measure of attractiveness of a zero-cost investment. The existence of high gain-loss ratio investments is related to SDFs exhibiting extreme deviations from m^*. Their assumption restricting good-deals is that excess payoffs should have a gain-loss ratio below an exogenous limit $\overline{L}$. This is equivalent to restricting the admissible set to SDFs that exhibit extreme deviations from m^* smaller than $\overline{L}$. The limit $\overline{L}$ must be chosen to exceed the maximum gain-loss ratio obtained with the basis assets $\overline{L}_B$. As $\overline{L} \to \infty$, their option price bounds converge to the no-arbitrage bounds. As $\overline{L}$ approaches $\overline{L}_B$, the bounds converge to the price implied by the SDF m' which is the smallest modification to m^*, according to the gain-loss ratio discrepancy measure, that correctly prices the basis assets. For a given $\overline{L}$, not all SDFs in their admissible set will be economically well-behaved, and it is necessary to choose $\overline{L}$ *a priori*.

In this context, our approach can again be viewed as a "well-behaved" good-deal pricing. While their price bounds are implied by measures that do not exceed an exogenous level $\overline{L}$ of gain-loss ratio discrepancy, we select only measures minimizing dispersion across a continuous family of divergences. That is, each measure we select is the smallest modification to the risk-neutral pricing kernel (or the empirical measure), according to a different Cressie-Read discrepancy, that correctly prices the basis assets. Even so, since each element in our admissible set is a minimum dispersion measure, all the remaining measures in the set have higher discrepancy from the perspective of a given γ. Therefore, our method imposes an endogenous maximum limit $\overline{L}$ to the Cressie-Read discrepancy of the risk-neutral measure and selects only measures consistent with HARA investors.

Stochastic Dominance

The stochastic dominance approach to option pricing was initiated by Perrakis and Ryan (1984), Levy (1985) and Ritchken (1985).[25] In order to derive bounds to option prices, this approach restricts the admissible set of SDFs that price the underlying asset and the risk-free bond to pricing kernels that are monotonically decreasing in the underlying returns. Therefore, the stochastic dominance approach identifies the range of values for a given option consistent with all risk averse investors, that is, any investor with increasing and concave utility function.

In the particular case of constructing option price bounds with the underlying asset as the single risky basis asset, the minimum dispersion admissible set will be contained in the set generated by the stochastic dominance method. This is because any minimum

[25] This method has been extended in a number of directions, such as to incorporate transactions costs (Constantinides and Perrakis, 2002).

dispersion measure is monotonically decreasing in the underlying returns. In fact, within the class of risk averse investors, we select investors maximizing HARA utility functions. An alternative interpretation is that our approach selects only the decreasing risk-neutral measures that minimize a Cressie-Read discrepancy with respect to the physical measure. Naturally, this implies that our approach leads to tighter bounds to option prices.

Throughout the paper, in order to assess the relevance of the more restrict information coming from the minimum dispersion admissible set, we compare the minimum dispersion price bounds with the stochastic dominance bounds. Besides providing tighter bounds that can potentially be more informative, our method can also be seen as a complement to the stochastic dominance approach, as it allows to identify which types of risk averse investors, from a considerably general class, price different options in the cross-section.

1.4 Pricing of Options in Simulated Economies

We start our analysis by exploring the implications of our method to option prices in simulated economies coming from well-known parametric models. With the simulations, we can isolate and understand the effects of different sources of risk and incompleteness on the relation between option prices and underlying returns. We calculate price bounds from underlying returns sampled from the physical distribution and compare them to option prices implied by the models under the risk-neutral parametrization. For options with different strikes and maturities, we also recover the implied γ, i.e., the marginal risk-neutral measure responsible for pricing each option.

The first economy considered is a Black-Scholes environment where the incompleteness comes from the absence of dynamic trading.[26] The second economy comes from the stochastic volatility and correlated jumps (SVCJ) model (Bates, 2000; Duffie, Pan and Singleton, 2000), incorporating additional sources of option market incompleteness. The SVCJ model captures important empirical stylized facts in equity markets and reliable estimates of its physical and risk-neutral parameters have been obtained in the option pricing literature. For instance, it is well documented for several markets that returns present stochastic volatility and jumps. Bates (2000) finds evidence against specifications of affine jump-diffusion models with pure stochastic volatility components or only jumps in returns and chooses as benchmark the SVCJ model. In Appendix B 1.8, we present in detail the models, the parametrization adopted and the simulation procedure.

[26] The Black-Scholes model is known to generate a complete market when the underlying stock and bond can be continuously traded to perfectly hedge an option payoff. However, in an environment without intermediate trading as we consider here, it generates an incomplete market and options will be non-redundant.

The simulation allows us to sample underlying returns directly from the true physical measure. To approximate the population value of the price bounds, we simulate one million underlying returns, compounded according to the option maturity, from the model-implied physical distribution. We calculate the minimum dispersion price bounds consistent with the set $[\underline{\gamma}, \overline{\gamma}]$ of strictly positive measures for European call options with several combinations of moneyness and maturity. Results for put options are analogous given put-call parity. For comparison, we calculate from the same returns the stochastic dominance bounds, as in Ritchken (1985).

1.4.1 Option Price Bounds

In a Black-Scholes economy, there is only one source of risk coming from the price process of the underlying asset. The market incompleteness is due to the absence of intermediate trading until the option expiration. We simulate lognormal returns from the physical distribution, calculate option price bounds and compare them with the option prices given by the Black-Scholes formula. Figure 1.4 plots the minimum dispersion option price bounds, the stochastic dominance (SSD) bounds and the Black-Scholes theoretical prices for call options with different combinations of time to maturity and moneyness. Prices are converted to implied volatilities for easier visualization. All option prices are contained in the minimum dispersion bounds, which in turn are contained in the wider SSD bounds. More specifically, the minimum dispersion price bounds are reasonably flat and close to the model constant implied volatility, while the SSD bounds get wider for options far-from-the-money. In this economy, option prices and underlying returns are perfectly consistent with each other.

The SVCJ model presents additional sources of risk that would introduce market incompleteness even on a continuous trading environment: stochastic volatility and jumps. This generates a more realistic economy, where log-returns under the physical distribution will no longer be Gaussian, being negatively skewed and leptokurtic instead. Figure 1.5 plots the option price bounds and the theoretical option prices implied by the SVCJ model converted to implied volatilities. The SVCJ model produces a volatility smile, where far-from-the-money options are relatively more expensive than near-the-money options, as represented by larger implied volatilities. The volatility smile is more pronounced for short maturity options, and it is evidence that the Black-Scholes model is not valid in this market. In particular, the smile reflects the fact that the tails of the SPD implied by the SVCJ model are thicker than the ones of a lognormal distribution.

For options with maturity greater than or equal to 3 months, all prices are contained in the minimum dispersion bounds, which are tighter than the stochastic dominance bounds, especially as options get far-from-the-money. The first panel of Figure 1.5 plots the bounds for 1-month options. In this case, OTM calls with moneyness (S/X) smaller than 0.93 violate the minimum dispersion upper bounds. These options also violate the SSD upper bounds, which are very close to the minimum dispersion ones in this moneyness region. As for the lower bounds, our approach improves a great deal on the SSD lower bounds, which reach zero prices for the OTM options. This suggests that the minimum dispersion bounds are able to capture the essential information for option prices from the class of risk averse investors.

The violations of 1-month OTM call prices happen because the SPD coming from the SVCJ model produces an implied volatility smile, while the underlying returns from the physical distribution generate bounds consistent with a volatility smirk. In other words, the model-implied option prices reflect a higher probability of large positive returns (making the OTM call finish ITM at expiration) than the minimum dispersion risk-neutral measures identified from the underlying returns. On the other hand, both the SVCJ model and HARA investors agree on the probability of large negative returns making OTM puts (and by put-call parity, ITM calls) pay at expiration.

The findings above can be summarized as follows. First, the minimum dispersion bounds seem to capture the essential information coming from risk averse investors to reconcile options and underlying returns. Second, option prices and underlying returns are perfectly consistent with each other when log-returns are Gaussian, even though the market is incomplete. With additional sources of risk and incompleteness given by stochastic volatility and jumps, there may exist violations of the bounds extracted from the returns. However, the vast majority of options considered under the SVCJ model are reconciled with underlying returns under the physical distribution.

1.4.2 Marginal Risk-Neutral Measures and Implied γ

In a Black-Scholes economy with no dynamic trading, Rubinstein (1976) shows that the Black-Scholes option price can be obtained by specifying a CRRA utility function. Given that the CRRA is a particular case of the HARA class, we should expect that there is a minimum dispersion risk-neutral measure that gives the correct option price of the Black-Scholes model. In the next proposition, we show that this is true and that the optimal discrepancy is completely characterized by the parameters of the Black-Scholes model.

Proposition 3. *Suppose a Black-Scholes economy with drift μ, risk-free rate r and volatility σ. Then, there is a Cressie-Read discrepancy indexed by γ^* for which the implied minimum dispersion option price equals the Black-Scholes price, given by:*

$$\gamma^* = -\frac{\sigma^2}{\mu - r} \; . \tag{1.17}$$

Proof. See Online Appendix 1.12. □

The proposition above states that a unique minimum dispersion risk-neutral measure should correctly price any option in a Black-Scholes economy, regardless of the moneyness and maturity. That is, there is a single HARA investor (in fact, a CRRA investor) that is marginal in the underlying asset, the risk-free bond and all options in the market at the same time. For the usual case that $\mu - r > 0$, we have that $\gamma^* < 0$ and the investor is always prudent. Smaller equity premia and higher volatilities imply that options require investors with larger absolute prudence to be priced. If the equity premium goes to zero, i.e., $\mu - r \to 0$, the optimal discrepancy is given by $\gamma^* \to -\infty$, which corresponds to the risk-neutral investor. This is consistent with the fact that when there is no risk premia in the economy, the physical measure already satisfies risk neutrality, implying that the minimum dispersion risk-neutral measure, according to any γ, is precisely $\mathbb{P}$. On the other hand, an arbitrarily large equity premium would be consistent with a CARA investor, as $\gamma^* \to 0$, which has the smallest prudence from all investors with $\gamma \leq 0$.

In other words, Proposition 3 states that the SPD implied by the minimum dispersion risk-neutral measure indexed by γ^* equals the lognormal SPD implied by the Black-Scholes model. This is confirmed in the left panel of Figure 1.6. Following Breeden and Litzenberger (1978), we recover both SPDs for maturity equal to three months.[27] As can be seen, the SPD implied by measure γ^* (equal to -0.8 given the parametrization considered) exactly matches the one of the model. Since this happens for all maturities, the implied γ surface, depicted in the right panel, is flat and approximately equal to γ^* for all combinations of moneyness and maturity. Two interpretations arise from these findings. First, the flat implied γ surface indicates the existence of a unique marginal investor in the market, that is, a representative investor. This confirms the results in Rubinstein (1976). Second, the constant implied γ means that options in the cross-section are equally expensive, in relative terms, given the underlying return physical distribution.[28] The same conclusion is achieved by looking at the constant implied volatility, which correctly assumes that underlying returns are lognormal.

When we consider the SVCJ model, underlying returns are no longer lognormal. In this case, the implied γ will take into account information from the true physical distribution, while the implied volatility will continue to denote deviations from Gaussian log-returns. This will naturally lead to different conclusions about relative option expensiveness. This is explained by the left panel of Figure 1.7, which depicts SPDs extracted

[27] That is, we calculate the SPD from the second derivative of the option prices implied by each model with respect to the strike. This method allows to compute in the same manner SPDs coming from different models. For the minimum dispersion risk-neutral measure, this is equivalent to a smoothed version of the histogram approximating the SPD, as in the right panel of Figure 1.2.

[28] We consider options on the same underlying asset to be equally expensive, in relative terms, when they are priced by the same risk-neutral measure coming from a given model.

from options with three months to maturity. We denote by SPD_{SVCJ}, $SPD_{ATM,BS}$ and $SPD_{ATM,\gamma}$ the SPD implied by the SVCJ model, by the Black-Scholes model with ATM implied volatility and by the minimum dispersion risk-neutral measure given by the ATM implied γ, respectively.

The $SPD_{ATM,\gamma}$ is a risk-neutral distribution minimizing dispersion with respect to the negatively skewed and leptokurtic physical distribution. As shown in Section 2, the minimum dispersion measure incorporates information from the physical distribution with a risk adjustment that results in a more negatively skewed distribution. In contrast, the $SPD_{ATM,BS}$ is lognormal, implying that the risk-neutral density of log-returns is symmetric. As can be seen in Figure 1.7, this is such that the $SPD_{ATM,\gamma}$ has a thicker left tail and a thinner right tail than the $SPD_{ATM,BS}$. This means that the minimum dispersion risk-neutral measure with ATM implied γ generates a volatility smirk, where ITM calls (OTM puts) are relatively more expensive than OTM calls (ITM puts) when measured by implied volatilities.[29]

If the SVCJ model generated the same volatility smirk as the ATM implied γ minimum dispersion risk-neutral measure, there would be a flat implied γ over strikes, even though the implied volatility is not flat. However, in this simulated economy, the SPD_{SVCJ} has thicker tails (both left and right) than the $SPD_{ATM,\gamma}$. This generates an implied γ smile for three-month to maturity options, as can be seen in the surface plotted in the right panel of Figure 1.7.[30] The implied γ smile indicates that when we consider information from the physical distribution, OTM calls (ITM puts) are relatively more expensive than ITM calls (OTM puts). In contrast, the volatility smile in the upper right panel of Figure 1.5 is almost a smirk, where ITM calls (OTM puts) are relatively more expensive than OTM calls (ITM puts). This difference arises because the required compensation with respect to the risk-neutral measure pricing the ATM option is smaller to price OTM puts than to price OTM calls under the implied γ, while the opposite is true under the implied volatility. That is, the left tail of the $SPD_{ATM,\gamma}$ is thicker than that of the $SPD_{ATM,BS}$ and closer to the left tail of the SPD_{SVCJ}, while the right tail of the $SPD_{ATM,BS}$ is thicker than that of the $SPD_{ATM,\gamma}$ and closer to the right tail of the SPD_{SVCJ}.

Most importantly, while the volatility smile reflects the misspecification of the Black-Scholes model, the implied γ smile can be economically interpreted as denoting the existence of heterogeneous marginal investors in a segmented option market. As can be

[29] This is because with a thicker left tail, large negative returns are more likely to happen than according to the lognormal distribution, increasing the probability that OTM puts finish ITM. On the other hand, a thinner right tail indicates a smaller probability of large positive returns that make the OTM call pay at expiration.

[30] Thicker tails with respect to $SPD_{ATM,\gamma}$ mean that OTM puts and OTM calls have higher probabilities of finishing ITM, therefore being relatively more expensive and requiring more negative γ's increasing the weights in extreme returns to be priced.

observed in the right panel of Figure 1.7, the heterogeneity of marginal investors decreases with the options time to maturity, leading to a flatter implied γ surface. That is, as time to maturity increases, $SPD_{ATM,\gamma}$ gets closer to SPD_{SVCJ}.

In sum, the implied γ represents a natural alternative to the implied volatility for drawing relative value comparisons between options in the cross-section that: (1) yields the same conclusions in a Black-Scholes economy, but (2) since it nonparametrically takes into account information from the true physical distribution of the underlying returns, yields different conclusions when returns are not lognormal, while (3) being economically interpretable and perfectly consistent with heterogeneous marginal investors in the market.

1.5 Pricing of S&P 500 Options

In this section, we apply our method to conduct a comprehensive analysis of the pricing of S&P 500 options given the underlying returns.[31] In particular, we use the minimum dispersion option price bounds to investigate whether option prices can be reconciled with returns on the index. In addition, we investigate how they are reconciled by recovering and interpreting the marginal risk-neutral measures from the cross-sections of options.

1.5.1 Data

We consider two sources of option data. For the period January 2, 1987 to December 29, 1995, we use the tick-by-tick Berkeley Options Database. We follow Jackwerth and Rubinstein (1996) in constructing a representative daily set of option prices from the raw data. The procedure is described in Appendix C 1.9. For the period January 4, 1996 to June 28, 2019, the data is obtained from OptionMetrics, consisting of end-of-day bid and ask quotes, volume, strike and expiration date for each option. As standard in the literature, we use the midpoint of the bid and ask quotes as the option prices and apply a handful of filters to the raw data. Observations with zero volume, bid-ask spread lower than 0.85, bid lower than 1/8 and implied volatility greater than 0.7 are dropped.

For both databases, we eliminate options that violate the usual no-arbitrage conditions. To calculate the time to expiration of each option, we take into account if the

[31] S&P 500 options are one of the most traded derivatives in the world, the most actively traded European options and have been the focus of many previous studies, such as Rubinstein (1994), Jackwerth and Rubinstein (1996), Ait-Sahalia and Lo (1998), Ait-Sahalia and Lo (2000), Jackwerth (2000), Ait-Sahalia, Wang and Yared (2001), Rosenberg and Engle (2002) and Constantinides, Jackwerth and Perrakis (2009).

contract settlement is at the open or close.[32] Moreover, since the S&P 500 index typically pays a dividend, we estimate the dividend yield from the put-call parity relation, using the pair of call and put options that are closest to ATM.[33] The option price bounds will be calculated from returns on the S&P 500 index and the risk-free rate. Daily data on the index is obtained from Bloomberg, for the sample covering from July 1, 1954 to June 28, 2019. We consider the 3-month Treasury Bill rate from the Saint Louis FRED database as a proxy for the risk-free rate, for the same period for which the option data is available.

We consider only options with time to expiration between 20 and 365 calendar days and moneyness between 0.90 and 1.10. The option data is divided in nine moneyness-maturity categories. With respect to moneyness, we group options in three intervals: OTM call (ITM put) if $S/X \in [0.90, 0.97)$, ATM if $S/X \in [0.97, 1.03)$ and ITM call (OTM put) if $S/X \in [1.03, 1.10]$. As for maturities, we classify options as short-term ($[20, 90)$ days), medium-term ($[90, 180)$ days) and long-term ($[180, 365]$ days). Our final sample from the Berkeley Database consists of 181,737 options, with 87,130 calls and 94,607 puts and an average of 80 options per trading day. From OptionMetrics, our final sample consists of 1,817,095 options, with 818,666 calls and 998,429 puts. There is an average of 307 options in the cross-section per trading day.

Summary statistics for OptionMetrics data after applying our filtering are reported in Table 1.1.[34] The average call price ranges from \$8.19 for short-term OTM to \$166.74 for long-term ITM, while the average put price is between \$13.49 for short-term OTM and \$145.96 for long-term ITM. Most calls are ATM, while most puts are OTM. Moreover, short-term options represent 78.49% of the total sample. Table 1.1 also reports the average implied volatilities in each moneyness-maturity category. The implied volatilities of call options are increasing in both moneyness and time to expiration. Put implied volatilities are also increasing in time to expiration, while decreasing in moneyness for long-term options. Short-term and medium-term puts present a U-shaped volatility pattern as the option goes from OTM to ATM and then to ITM, where the OTM put implied volatilities take the highest values. The differences between implied volatilities in the cross-section are more pronounced for short-term options.

---

[32] For PM settled options, the time to expiration is the number of days between the trade date and the expiration date. For AM settled options, we use the number of days between the dates less one.

[33] For each day in the sample and each maturity T, we identify the pair of ATM call and put option prices, c and p, and estimate the dividend yield as $q = -(1/T)\ln[(c - p + Xe^{-rT})/S]$, where r is the corresponding continuously compounded risk-free rate. This procedure is usually preferred to the backward-looking approach of estimating the future dividend rate using past daily dividend payments on the index.

[34] We report the summary statistics for the Berkeley Database in the Online Appendix 1.12.

1.5.2 Estimation

To calculate the minimum dispersion option price bounds, we need to estimate the physical distribution of the underlying returns. We proceed by estimating at each point in time the conditional return distribution without assuming a parametric form. This is done in two sequential steps. First, for each date t in our option sample, and for each time to maturity T in the option cross-section at date t, we estimate the unconditional return distribution as the histogram of overlapping T-day returns calculated using the index sample before date t. We impose the economic restriction of a 5% lower bound on the annualized equity premium over the risk-free rate. That is, if for a date t and maturity T the annualized mean of the unconditional return distribution generates a premium less than 5% over the risk-free rate, we demean the returns and reintroduce a 5% equity premium.[35]

Second, we make this distribution conditional by adjusting its volatility at each time t. We avoid backward-looking procedures such as historical volatility or potentially misspecified parametric models such as GARCH and realized variance models. Instead, we estimate the conditional physical volatility by discounting a premium from the implied volatility of the ATM call option with maturity T. It is well-known that the risk-neutral index volatility, usually measured by the ATM implied volatility, generally exceeds the physical return volatility.[36] We use the simulations of the SVCJ model in Section 1.4 to estimate the premium for different maturities.[37] Then, for each date t and time to maturity T, we scale the unconditional distribution to have standard deviation equal to our conditional forward-looking estimate of physical volatility.

The calculation of the bounds also use the index current price, the option strike and the risk-free rate at date t compounded to the appropriate maturity. To account for the dividends, we discount the current price S to Se^{-qT} and we add the dividend yield with corresponding maturity to the mean of the conditional return distribution. From this distribution, where each T-day return is a realization of a state of nature, we calculate the minimum dispersion price bounds. From the same returns, we also calculate the stochastic dominance (SSD) bounds.

The stochastic dominance approach allows for pricing kernels in the restricted admissible set that have zero values in some states of nature. To account for this possibility,

[35]Despite the fact that results are practically unchanged by this restriction, we keep it because it is economically sound to consider a lower bound to the equity premium. See, for instance, Martin (2017).

[36]See Bakshi and Madan (2006) and the references therein.

[37] For short-term options, we estimate the premium as the difference, in the SVCJ model with 1 month to maturity, between the model ATM implied volatility and the annualized volatility of the one million returns drawn from the physical distribution. For medium- and long-term options, we do the same but with the SVCJ model with 3 and 6 months to maturity, respectively. The estimated annualized short-, medium- and long-term premia are 0.38%, 0.76% and 1.11%, respectively.

in addition to the bounds implied by the strictly positive minimum dispersion risk-neutral measures, we also calculate a version of our bounds that allows for measures with weights equal to zero in some states of nature. We do that by allowing for greater γ's in the right hand side of the strictly positive admissible set $[\underline{\gamma}, \overline{\gamma}]$. In particular, we set $\overline{\gamma}' = 40$.[38] Given that option prices are decreasing in γ, this generates a smaller lower price bound, while maintaining the same upper bound. From now on, we shall refer to these bounds as the minimum dispersion (MD) bounds, and to the bounds selecting only strictly positive measures as minimum dispersion no-arbitrage (MDNA) bounds. When a given option price is observed to be inside the calculated MD bounds, we also compute the implied γ that correctly prices the option.

1.5.3 Empirical Results

Option Price Bounds

We calculate price bounds from the underlying returns for each option in our sample. Table 1.2 reports, for the OptionMetrics data, the percentage of times that observed option prices are contained in the bounds, the percentage of upper and lower bound violations and the average tightness of the bounds, for each option category.[39] Focusing first on the aggregate results, nearly all option prices (99.42% of the calls and 98.63% of the puts) are consistent with the SSD bounds, i.e., with risk averse investors. Importantly, almost all of the observed prices are also contained in the tighter MD bounds (98.02% of the calls and 96.72% of the puts), confirming that the minimum dispersion admissible set identifies a comprehensive class of risk averse investors providing relevant information about option prices.

Detailing the analysis at the level of option categories, the options that most violate the bounds are short-term ITM calls and ITM puts, where 94.25% and 88.65% of these options are contained by the MD bounds, respectively. This may be due to the fact that ITM options are less liquid and may present some unreliable prices. On the other hand, 97.53% of the short-term OTM puts are consistent with the MD bounds. This contradicts the common notion that the left tail of the risk-neutral distribution is hard to be reconciled. As the time to maturity increases for medium- and long-term options, the percentage of options contained in the bounds increases. That is, the longer the maturity, the easier it is to reconcile the information from option prices and underlying returns.

We also calculate the MDNA bounds, identifying prices consistent with strictly

[38] We choose $\overline{\gamma}' = 40$ as this is a large γ that will always be above the $\overline{\gamma}$ defining the strictly positive admissible set and for which the implied option price will be close to the SSD lower bound.

[39] We report the results for the Berkeley Database in the Online Appendix 1.12. Results are similar and provide additional robustness to our analysis.

positive risk-neutral measures. The MDNA bounds capture 71.23% of the calls and 87.46% of the puts, leaving unexplained a non-negligible portion of the prices. However, a specific pattern appears when we look at the categories of options. ITM calls and OTM puts can be explained by strictly positive measures, with the much tighter MDNA bounds capturing most of the observed prices. In contrast, OTM calls and ITM puts clearly require risk-neutral measures with zeros in some states of nature to be priced, as the number of MDNA lower bound violations is large. As we show in Section 1.3, the measures associated to γ's above $\overline{\gamma}$ set to zero states corresponding to the largest positive returns. Therefore, to reconcile the right tail of the risk-neutral distribution, risk-neutral measures identified from the physical distribution of underlying returns need to decrease the probability mass in large positive returns to make them compatible with option prices.

Overall, as the vast majority of option prices lie within the MD and SSD bounds, option prices are mostly consistent with underlying returns when we entertain the possibility that markets are incomplete. The tighter MDNA bounds contain most of the prices of OTM puts (and ITM calls), while some of the ATM options and most of the OTM calls (and ITM puts) need the wider MD bounds to be reconciled. More specifically, ATM and high-strike options require smaller lower bounds consistent with risk-neutral measures setting to zero the largest positive returns.

Marginal Risk-Neutral Measures and Implied γ

Having established that option prices can be reconciled with returns on the underlying asset, we now use our method to analyze in detail how they are reconciled. Table 1.3 reports the average implied γ over the 1996-2019 sample for each option category.[40] On average, the implied γ of call options is decreasing in moneyness (S/X), while it is increasing for put options. That is, option prices present a smirk of the implied γ, indicating the existence of heterogeneous marginal investors in a segmented option market. OTM puts (and ITM calls) are priced by investors with positive prudence, convex marginal utilities and aversion to downside risk. This is consistent with the common notion that OTM puts on the S&P 500 index are demanded as crash-insurance instruments (Rubinstein, 1994). On the other hand, OTM calls (and ITM puts) require large positive γ's to be priced, associated to concave marginal utilities. Since these investors have negative prudence, this result is in agreement with the demand of OTM calls by speculators buying them as a leveraged bet. The heterogeneity of marginal investors is more pronounced for short-term options, decreasing as the maturity increases.

To assess how the implied γ patterns change over time, the upper pannels of Fig-

[40] We report the average implied γ over the 1987-1995 sample in the Online Appendix 1.12. The results are qualitatively similar.

ure 1.8 plot the two-month moving average of the implied γ and the $\underline{\gamma}$ and $\overline{\gamma}$ defining the MDNA bounds, for short-term calls and puts with different moneyness.[41] The lower panels plot the corresponding moving averages of the implied volatilities. There is little variation in the implied γ of OTM puts and ITM calls, which are almost always reconciled by negative γ's associated to prudent investors. ATM options also do not display much variation, being in great part consistent with strictly positive risk-neutral measures inside the interval $[\underline{\gamma}, \overline{\gamma}]$. In contrast, the implied γ of OTM calls and ITM puts varies considerably over the sample. While these options are mainly reconciled by large positive γ's, there are periods where the implied γ becomes negative. In other words, the heterogeneity of marginal investors is time-varying and is mainly driven by the implied γ of OTM calls and ITM puts. In particular, the heterogeneity tends to decrease with financial crises. This is more evident for the 2008 crisis, where the marginal investors of OTM calls and ITM puts switch from non-prudent to prudent behavior in face of the market downturn.

The implied γ also represents an alternative way of interpreting the structure of option prices, i.e., of drawing comparisons between prices of options in the cross-section. In contrast to the implied volatility, these comparisons directly take into account information from the physical distribution of underlying returns. The ATM implied γ SPD represents the "closest" distribution to the physical distribution that is risk-neutral and consistent with the level of option prices.[42] The implied γ smirk means that, with respect to this SPD, more negative γ's that put more weight in large negative returns (that generate an SPD with thicker left tail) are necessary to price OTM puts, while larger γ's that put less weight in large positive returns (that generate an SPD with thinner right tail) are necessary to price OTM calls. The extent to which these compensations are necessary in the cross-section vary through time, providing new insights into the relative expensiveness and skew patterns of options, which we discuss below.

First, under the implied γ, for most dates the difference in relative expensiveness between low-strike and ATM options is small, while that between high-strike and ATM options is large. The opposite is true according to the implied volatility. In other words, the left tail of the ATM implied γ SPD is usually thinner but close to that of the SPD correctly pricing all options in the cross-section (i.e., the option-implied SPD), while the right tail is thicker. In contrast, the left tail of the ATM implied volatility lognormal SPD is always much thinner than that of the option-implied SPD. This indicates that, when we consider information from the underlying return physical distribution, OTM puts (and ITM calls) are generally not too expensive.

[41] That is, first, for each day, we average the implied γ over options in each moneyness and maturity category. Then, we calculate the two-month moving average of the daily mean implied γ's in each category over our sample. The plots for medium and long-term options are similar and are available in the Online Appendix 1.12.

[42] Closest in the sense of risk-neutral measures that minimize dispersions in the Cressie-Read family.

An alternative way of illustrating the differences above is by comparing the implied volatility curves generated by the observed option prices and the price bounds. Figure 1.9 plots, in the upper left panel, price bounds and option prices with two months to maturity converted to implied volatilities for the date January 28, 2015. We choose this date as it is representative of what we find in most of the dates in the sample. The upper right panel depicts the corresponding implied γ curve. The MDNA bounds contain the ATM and low-strike option prices. Visually, the volatility smirks generated by the MDNA upper and lower bounds display a similar slope to the option-implied volatility smirk in the moneyness region $S/X \in [1.0, 1.10]$. Thus, we should expect that little compensation is necessary to price the low-strike options. This is confirmed by the nearly flat implied γ for these strikes. On the other hand, the volatility smirk for the high-strike options is more downward sloped, being only contained by the wider MD and SSD bounds. This implies that, as the moneyness decreases, these options require more positive γ's to be priced, as shown in the upper right panel.

A second important pattern that can be identified is that there are periods where the implied γ is approximately flat, even though there is a smirk of the implied volatility. This is more evident in the 2008 financial crisis.[43] In these periods, the ATM implied γ SPD gets close to the option-implied SPD, mainly because the implied γ of high-strike options decreases. That is, OTM calls get relatively more expensive than before as they require smaller γ's (approaching the ATM implied γ) that increase the weights in large positive returns, because investors believe that upward movements in the market are more likely to happen. On the other hand, the ATM implied volatility lognormal SPD continues to have a thicker right tail than the option-implied SPD. To better illustrate, the lower panels of Figure 1.9 report the price bounds, option prices and implied γ with two months to maturity converted to implied volatilities for December 19, 2008. This date is representative of what happens during the 2008 financial crisis. As can be observed, all option prices are contained by the MDNA bounds, where the slope of the option-implied volatility smirk is very similar to the slopes of the volatility smirks generated by the bounds. This is such that there is a nearly flat implied γ across all strikes, indicating that there is approximately a single, prudent marginal investor in the whole cross-section.

In summary, option prices are reconciled by heterogeneous risk-neutral measures depending on the option moneyness and maturity. OTM puts are priced by prudent investors averse to market downside risk, while investors pricing OTM calls most of the time are nonprudent. As the option maturity increases, the heterogeneity in the cross-section decreases. Moreover, the heterogeneity represents an alternative way of assessing the relative expensiveness and skew patterns of options with different strikes and maturities.

[43] See also in the end of 1997, in the end of 2003 and in the beginning of 2005 for calls and puts, and in the end of 2000 for puts.

This provides new insights into the pricing of options given the information incorporated from underlying returns.

Pricing of S&P 500 Options Before and After the 1987 Crash

The crash of the stock market in October 1987 is known to have changed the patterns of implied volatility in S&P 500 options. Rubinstein (1994) shows that prior to the crash, the index implied volatilities were approximately flat, while afterwards they usually display a smirk: the volatility increases as the strike price decreases. That is, OTM puts (and ITM calls) are more expensive than ITM puts (and OTM calls) in terms of implied volatility. In related work, Jackwerth and Rubinstein (1996) recover the risk-neutral distributions from option prices before and after the market crash, identifying a significant change in shape between them. While the precrash distributions resemble the lognormal distribution, the postcrash ones have thicker left tails and thinner right tails.[44]

In this context, we use our method to provide novel insights on how the structure of option prices has changed from before to after the crash, given the underlying returns. To that end, the upper pannels of Figure 1.10 plot, for the period 1987-1995, the two-month moving average of the implied γ and the $\underline{\gamma}$ and $\overline{\gamma}$ defining the MDNA bounds, for short-term calls and puts with different moneyness.[45] The lower panels plot the corresponding moving averages of the implied volatilities. The vertical dashed line in each plot marks the 1987 crash. Focusing first on the implied volatilities, we can see that prior to the crash they are very close to each other in the cross-section. After the crash, the usual pattern of the volatility smirk appears, where ITM calls and OTM puts are relatively more expensive than ATM and high-strike options. Rubinstein (1994) suggests that the postcrash volatility smirk is due to "crashophobia": investors are concerned about the possibility of another crash, making OTM puts (and ITM calls) to become more highly valued.

Focusing now on the implied γ, a strikingly different pattern can be identified. OTM puts (and ITM calls) were actually cheaper (with large positive implied γ's), in relative terms, than ATM and high-strike options prior to the crash. That is, the ATM implied γ SPD had a thicker left tail than the option-implied SPD, which was close to lognormal. This indicates that investors were pricing options using the Black-Scholes formula, even though this was inconsistent with the physical distribution. In fact, from Section 4, if the physical distribution were lognormal, the ATM implied γ SPD and the implied γ curve would have been lognormal and flat, respectively. The mismatch between

[44] See also Ait-Sahalia and Lo (1998) on the comparison between the postcrash option-implied SPD and the ATM implied volatility lognormal SPD.

[45] The plots for medium and long-term options are similar and are available in the Online Appendix 1.12.

the implications of underlying returns for option prices and how investors were pricing options is also evident in Figure 1.11. The left panel plots option price bounds and option prices with two months to maturity converted to implied volatilities for the date May 26, 1987, while the right panel depicts the corresponding implied γ curve. This date is representative of what we find in most of the dates prior to the crash. As can be seen, the implied volatility is flat across strikes, while the MDNA, MD and SSD bounds generate volatility smirks. This mismatch produces the reverse implied γ smirk on the right panel, where OTM calls (and ITM puts) are relatively more expensive than ITM calls (and OTM puts).

After the crash, the structure of option prices changes to the implied γ smirk discussed in the previous subsection. OTM puts (and ITM calls) become relatively more expensive, as the option-implied SPD assigns a higher probability to large negative returns than the ATM implied γ SPD. This confirms the "crashophobia" phenomenon when we consider information from the physical distribution, but to a lesser extent than suggested by the implied volatility smirk, as for several periods the implied γ's of low-strike and ATM options are close to each other. On the other hand, OTM calls (and ITM puts) become relatively cheaper, where the right tail of the option-implied SPD is now thinner than that of the ATM implied γ SPD. In addition, the implied γ gets approximately flat at some periods (as after the 1991 crisis), even though there is a smirk of the implied volatility.

1.6 Conclusion

We introduce a new method to reconcile option prices with underlying returns in incomplete markets. From the returns on the underlying asset, we construct meaningful option price bounds and show how to recover the implied γ, a parameter uniquely identifying the marginal investor pricing a given option. The implied γ represents an alternative to the implied volatility for drawing relative value comparisons between options in the cross-section, that directly takes into account information from the physical distribution of underlying returns.

We apply our method to nearly two million S&P 500 options, finding that option prices are mostly consistent with underlying returns, as the vast majority of prices lie within the bounds we calculate. Option prices are reconciled by heterogeneous risk-neutral measures that depend on the option moneyness and maturity, indicating the existence of marginal investors with different assessments of market downside risk. The heterogeneity is time-varying and decreases during financial crises. Furthermore, the implied γ patterns provide novel insights into the relative expensiveness and skew patterns of index option prices before and after the 1987 crash.

Our paper suggests several possible avenues for future research. First, it has been documented that there is a significant difference between the volatility smiles of index and individual stock options, despite a relative similarity of their underlying returns (Bakshi, Kapadia and Madan, 2003). Applying our method to individual stock options would provide a formal investigation of how different their structure of prices is from that of index options, directly taking into account information from the physical distributions. Second, while we show with simulations that stochastic volatility and jumps can generate an implied γ smile, the implied γ patterns in our empirical application are likely to be also explained by demand-pressure effects as in Gârleanu, Pedersen and Poteshman (2009). It would be interesting to analyze the extent to which the relative expensiveness and skew patterns identified by the implied γ can be explained by the variation in demand of index options, especially when these patterns differ from those of implied volatility.

While we focus on reconciling the cross-section of S&P 500 option prices with the time series of underlying returns, our method can also be applied to price options in illiquid markets. The price bounds we provide can be used to identify a range of reasonable values for options given the information from the underlying returns. More than that, our analysis connects the prices of options with different strikes and maturities to specific risk-neutral measures associated to different risk preferences. This could help traders or market-makers with views on the preferences of marginal investors in obtaining economically sound option prices and building a trading history in the illiquid market.

In addition, our method can be extended to hedge options. For a given option to be hedged, one could identify the implied γ and obtain the sensitivity of the option price to perturbations in traditional risk factors, such as the underlying asset price, volatility and risk-free rate. Keeping fixed the implied γ for that option, one could re-estimate the risk-neutral measure following each perturbation, obtaining new option prices and the corresponding "greeks" that can be used to hedge the option.

1.7 Appendix A - Additional Interpretations of Minimum Dispersion Risk-Neutral Measures

1.7.1 Bayesian Interpretation

Problem (1.4) has a Bayesian interpretation. The physical measure $\mathbb{P}$ represents a prior to the risk-neutral measure $\mathbb{Q}$, while the restrictions in (1.4) define risk neutrality. With data on basis assets returns available, these restrictions can be imposed and used to update the prior according to a measure of information gain. In this sense, the minimum dispersion risk-neutral measures are posterior probabilities that add the minimum amount of information needed for $\mathbb{P}$ to price the basis assets, taking into account different sensitivities to higher moments of the state-price density. To show that, we Taylor expand the expected value of $\phi_\gamma(\mathrm{d}\mathbb{Q}/\mathrm{d}\mathbb{P}) = \frac{(\mathrm{d}\mathbb{Q}/\mathrm{d}\mathbb{P})^{\gamma+1}-1}{\gamma(\gamma+1)}$ around 1. Note that $\phi_\gamma(1) = 0$, $\phi'_\gamma(x) = \frac{x^\gamma}{\gamma}$, $\phi''_\gamma(x) = x^{\gamma-1}$, $\phi'''_\gamma(x) = (\gamma-1)x^{\gamma-2}$, $\phi''''_\gamma(x) = (\gamma-1)(\gamma-2)x^{\gamma-3}$ and so on. Taylor expanding ϕ_γ and taking expectations on both sides, we obtain:

$$\mathbb{E}\left(\phi_\gamma\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)\right) = \frac{1}{2}\mathbb{E}\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} - 1\right)^2 + \frac{(\gamma-1)}{3!}\mathbb{E}\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} - 1\right)^3 + \frac{(\gamma-1)(\gamma-2)}{4!}\mathbb{E}\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} - 1\right)^4 + \ldots \tag{1.18}$$

Figure 1.12 depicts the weights given by different discrepancies to higher-order moments of the risk-neutral probabilities. As can be seen in the left panel, all discrepancies give the same weight to the variance. The Euclidean divergence ($\gamma = 1$) is a quadratic loss function, giving zero weight to all higher moments. Discrepancies with γ close to one, as the KLIC ($\gamma = 0$), give small weights to higher moments as compared to the variance. The Pearson's Chi-Square divergence ($\gamma = -2$) gives equal (absolute) weight to all moments of the state-price density, while more negative γ's increasingly weight more, in absolute terms, higher moments. The right panel of Figure 1.12 illustrates better this relation. For γ's greater than -2, the absolute weight given to risk-neutral measure moments is decreasing with the moment order. In contrast, $\gamma = -2$ gives equal absolute weights to all moments, while for $\gamma < -2$ the absolute weight increases with the moment order. Moreover, for all discrepancies, weights for even moments are positive and for odd moments are negative.

1.7.2 Portfolio Interpretation

The seminal work of Hansen and Jagannathan (1991) initiated a literature that proposes moment restrictions to pricing kernels in order to diagnose asset pricing models. They derive a minimum variance SDF that places a lower bound on the variance of candidate pricing kernels. This approach is generalized by Almeida and Garcia (2017) by deriving the Cressie-Read family of minimum discrepancy SDFs. Basically all SDF moment restrictions proposed in the literature are particular cases of the Cressie-Read family. The second moment restriction of Hansen and Jagannathan (1991) obtains for $\gamma = 1$ (for the nonnegative SDF case). Snow (1991) proposes restrictions on the p-th moment of the pricing kernel, with $p > 1$, which corresponds to $\gamma > 0$. Stutzer (1995) derives an information bound minimizing the KLIC ($\gamma = 0$), while Bansal and Lehmann (1997) obtain a growth-optimal bound that minimizes entropy ($\gamma = -1$). More recently, Liu (2020) complements the method of Snow (1991) by restricting the p-th moment of the SDF, with $p \in (-\infty, 1)$, corresponding to $\gamma < 0$.[46]

Hansen and Jagannathan (1991) show that the SDF volatility restriction has a portfolio interpretation: it is equivalent to an upper bound limit on the Sharpe ratio. In this section, we show that a restriction on the Cressie-Read dispersion of a pricing kernel also has a portfolio interpretation. It is equivalent to an upper bound on the generalized Sharpe ratio (Cerny, 2003) defined by the corresponding HARA utility function. We follow Cerny (2003) in measuring the attractiveness of a self-financing investment by the certainty equivalent of the resulting wealth W relative to the wealth of a riskless investment. The value of the best deal in the market with excess returns $\mathbf{R}$, denoted $\alpha(\mathbf{R})$, is defined implicitly as:

$$u^{\gamma}(W_0 R_f + \alpha(\mathbf{R})) \equiv \max_{\lambda} \mathbb{E}[u^{\gamma}(W_0 R_f + \lambda'(\mathbf{R} - \mathbf{R}_f))]. \tag{1.19}$$

The next proposition derives the SDF moment restrictions implied by the class of HARA utility functions.

Proposition 4. *Consider the class of HARA utility functions as in (1.12) indexed by γ. Then, the following SDF moment restrictions hold for $\gamma \in (-\infty, \infty)$:*[47]

$$\left(1 - \gamma A^{\gamma}(w_0)\alpha_{\gamma basis}\right)^{-(\gamma+1)} \leq \mathbb{E}[m^{\gamma+1}] \leq \left(1 - \gamma A^{\gamma}(w_0)\overline{\alpha}_{\gamma}\right)^{-(\gamma+1)}, \tag{1.20}$$

$$s.t. \;\; \mathbb{E}[m(\boldsymbol{R} - \boldsymbol{R}_f)] = \boldsymbol{0}, \;\; m \geq 0,$$

where $w_0 = W_0 R_f$, $A^{\gamma}(w_0) = a/(b - \gamma a w_0)$, $\overline{\alpha}_{\gamma} \geq \alpha_{\gamma basis}$ and $\alpha_{\gamma basis}$, the certainty equivalent of the best deal attainable in the market containing only the basis assets, is given by:

$$\alpha_{\gamma basis} = \min_{m} \;\; \alpha_{\gamma}(m), \;\; s.t. \;\; \mathbb{E}[m(\boldsymbol{R} - \boldsymbol{R}_f)] = \boldsymbol{0}, \;\; m \geq 0. \tag{1.21}$$

The restrictions above can be interpreted as reward-for-risk measures as follows:

$$\left(1 + h^2_{\gamma basis}\right)^{\frac{\gamma(\gamma+1)}{2}} \leq \mathbb{E}[m^{\gamma+1}] \leq \left(1 + \overline{h}^2_{\gamma}\right)^{\frac{\gamma(\gamma+1)}{2}}, \tag{1.22}$$

$$s.t. \;\; \mathbb{E}[m(\boldsymbol{R} - \boldsymbol{R}_f)] = \boldsymbol{0}, \;\; m \geq 0,$$

where $h^2_{\gamma} = 2A^{\gamma}(w_0)\alpha_{\gamma}$ is the generalized Sharpe ratio.

Proof. See Online Appendix 1.12. □

For a given γ, the left hand side restriction in (1.22) implies that minimizing the SDF Cressie-Read discrepancy is equivalent to maximizing the generalized Sharpe

[46] Except for the Hansen and Jagannathan (1991) case without the nonnegativity constraint, all these restrictions imply nonnegative pricing kernels, so we can talk interchangeably about SDFs and their risk-neutral measure counterparts.

[47] The limiting cases of $\gamma \to 0$ and $\gamma \to -1$ are proved in Cerny (2003).

ratio in the market with basis assets. This provides a portfolio interpretation for the minimum dispersion risk-neutral measures and the SDF moment restrictions that have been employed to diagnose asset pricing models. In contrast to the standard Sharpe ratio (that is associated to quadratic utility), generalized Sharpe ratios provide a consistent ranking of investment opportunities when asset returns are nonnormal (see Cerny, 2003).

Furthermore, suppose we were interested in finding the set of prices of a non-redundant asset that does not provide deals better than $\overline{\alpha}_\gamma$. Then, the right hand side restrictions imply that the extended market does not provide deals better than $\overline{\alpha}_\gamma$, as measured by the HARA utility indexed by γ, if and only if the non-redundant asset is priced with SDFs satisfying (1.20) (or (1.22)). This generalizes the equilibrium restrictions of Cochrane and Saa-Requejo (2000), as limiting the maximum SDF dispersion allowed in the admissible set is equivalent to restricting the maximum generalized Sharpe ratio attainable in the extended market. Their restrictions obtain for $\gamma = 1$.

1.8 Appendix B - Parametric Models for the Underlying Asset Dynamics

The parametric models for the underlying asset dynamics considered in the simulations allow for analytic solutions to the option pricing equation (1.15) under the models risk-neutral distributions. They also allow us to sample underlying returns directly from the model-implied physical distribution.

1.8.1 Black-Scholes Model

In the Black-Scholes model, the only source of risk is given by the underlying price process. The underlying asset price follows a geometric Brownian motion:

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t, \tag{1.23}$$

where S_t is the stock price at time t, μ is the average stock return, σ is the constant instantaneous volatility of the process and W_t is a standard Wiener process. The corresponding stochastic differential equation (SDE) under the risk-neutral measure is:

$$\frac{dS_t}{S_t} = r dt + \sigma dW_t(\mathbb{Q}), \tag{1.24}$$

where r is the risk-free rate. We assign standard values for the parameters: the drift is $\mu = 10\%$, the volatility $\sigma = 20\%$ and the risk-free rate $r = 5\%$.

The Black-Scholes model implies a lognormal distribution for stock returns under the physical measure $\mathbb{P}$:

$$\ln(R_k) \sim N((\mu - \sigma^2/2)T, \sigma^2 T), \tag{1.25}$$

from which we can draw returns and construct the distribution of future possible prices for the underlying asset for each time to expiration.

1.8.2 Stochastic Volatility and Correlated Jumps Model

The SVCJ model is defined by two SDEs respectively for the price and volatility of the underlying asset. The jumps in the price and volatility process happen at the same time and are perfectly correlated. The stock price, S_t, and its spot variance, V_t, solve:

$$dS_t = S_t\left(\mu - \overline{\mu}_s\lambda\right)dt + S_t\sqrt{V_t}dW_t^s + d\left(\sum_{n=1}^{N_t} S_{\tau_n-1}[e^{Z_n^s} - 1]\right), \tag{1.26}$$

$$dV_t = \kappa_v\left(\theta_v - V_t\right)dt + \sigma_v\sqrt{V_t}dW_t^v + d\left(\sum_{n=1}^{N_t} Z_n^v\right), \tag{1.27}$$

$$\overline{\mu}_s = \exp\left(\mu_s + \sigma_s^2/2\right) - 1, \tag{1.28}$$

where W_t^s and W_t^v are Brownian motions with correlation ρ; μ is the rate of return of the asset; κ_v is the rate at which the instantaneous variance V_t reverts to the long-run variance θ_v; σ_v is the volatility of the volatility generating process; N_t is the number of jumps until time t described as a Poisson process with intensity λ; $Z_n^s|Z_n^v \sim N(\mu_s + \rho_s Z_n^v, \sigma_s^2)$ are the jumps in prices; $Z_n^v \sim \exp(\mu_v)$ are the jumps in volatility and τ_n is the time of the n^{th} jump.

The market generated by the SVCJ model is incomplete even with dynamic trading, implying the existence of an infinity of theoretical risk-neutral measures consistent with the prices of the underlying asset. In particular, when considering jumps, the Girsanov theorem imposes very weak conditions for the jump distribution change of measure. The usual procedure in this context is to parametrize the possible changes of measure and make specific assumptions about the distributions of jumps. We follow Duffie, Pan and Singleton (2000) and Broadie, Chernov and Johannes (2007) in considering the following SDEs under the risk-neutral measure where options are priced:

$$dS_t = S_t\left(r - \overline{\mu}_s^{\mathbb{Q}}\lambda^{\mathbb{Q}}\right)dt + S_t\sqrt{V_t}dW_t^s(\mathbb{Q}) + d\left(\sum_{n=1}^{N_t(\mathbb{Q})} S_{\tau_n-1}[e^{Z_n^s(\mathbb{Q})} - 1]\right), \tag{1.29}$$

$$dV_t = \kappa_v^{\mathbb{Q}} \left(\theta_v^{\mathbb{Q}} - V_t \right) dt + \sigma_v \sqrt{V_t} dW_t^v(\mathbb{Q}) + d \left(\sum_{n=1}^{N_t(\mathbb{Q})} Z_n^v(\mathbb{Q}) \right), \tag{1.30}$$

$$\overline{\mu}_s^{\mathbb{Q}} = \exp\left(\mu_s^{\mathbb{Q}} + (\sigma_s^{\mathbb{Q}})^2/2 \right) - 1. \tag{1.31}$$

The absolute continuity requirement implies that some parameters (or combinations of parameters) should be the same under the physical and risk-neutral measures. In our case, this is true for σ_v, ρ and the product $\kappa_v \theta_v$. We follow Broadie, Chernov and Johannes (2007) in assuming that the intensity λ is a constant, that $Z_n^s(\mathbb{Q}) \sim N(\mu_s^{\mathbb{Q}}, (\sigma_s^{\mathbb{Q}})^2)$ and $Z_n^v(\mathbb{Q}) \sim \exp(\mu_v^{\mathbb{Q}})$. Furthermore, since the literature finds insignificant estimates for ρ_s, they also consider that $\rho_s = 0$. We take as a reference the physical parameters in Eraker, Johannes and Polson (2003), estimated using S&P 500 returns, and the risk-neutral parameters in Broadie, Chernov and Johannes (2007), estimated using S&P 500 futures option prices. The risk-free rate r is set to 5% and we consider the initial variance equal to its long-run mean $\theta_v + (\mu_v \lambda_v)/\kappa$. Table 1.4 reports the parameters.

In contrast to the Black-Scholes model, in the SVCJ model we must approximate the continuous time process by a discrete time one in order to simulate the data, which in turn might generate discretization bias. The conventional way to generate data from stochastic volatility models is to use Euler discretization. However, Broadie and Kaya (2006) find that the bias may be very large in some cases even if a large number of steps are used. They describe a method to simulate data from the exact distribution, effectively reducing discretization bias. For this reason, we work with exact sampling from the SVCJ model by using the method described by Broadie and Kaya (2006).

1.9 Appendix C - Berkeley Options Database

We obtain from the Berkeley Database all reported quotes covering S&P 500 options traded on the Chicago Board Options Exchange from January 2, 1987 to December 29, 1995. We conduct a similar procedure as Jackwerth and Rubinstein (1994) to construct from the tick-by-tick data a representative daily set of option prices. We drop the first and last 30 minutes of quotes each day and all quotes with bid-ask spread lower than 0.85. We also only keep options with moneyness between 0.90 and 1.10. From the remaining bid and ask quotes, we calculate option prices as the midpoint and translate each price to the corresponding Black-Scholes implied volatility. Then, for each day in our sample, we calculate the average implied volatility for each strike price and time to expiration. Finally, we translate the average implied volatilities back into option prices. This constitutes our representative daily set of option prices.

1.10 Tables

Table 1.1: **Summary Statistics of S&P 500 Index Options (1996-2019)**

	Call Options				Put Options			
Moneyness	Short	Medium	Long	Subtotal	Short	Medium	Long	Subtotal
	\$8.19	\$19.49	\$42.89	\$15.31	\$106.60	\$122.70	\$145.96	\$114.78
[0.90, 0.97)	13.11%	13.58%	15.26%	13.51%	16.35%	16.50%	17.01%	16.47%
	{162148}	{50745}	{34214}	{247107}	{53070}	{12952}	{10645}	{76667}
	\$34.96	\$62.39	\$97.08	\$42.10	\$38.46	\$66.71	\$95.32	\$45.73
[0.97, 1.03)	13.17%	15.62%	17.50%	13.73%	14.02%	16.19%	17.65%	14.53%
	{373929}	{51221}	{29643}	{454793}	{364781}	{53293}	{30968}	{449042}
	\$113.01	\$137.13	\$166.74	\$120.89	\$13.49	\$36.34	\$63.45	\$19.73
[1.03, 1.10]	18.83%	19.05%	19.60%	18.93%	18.43%	18.86%	19.48%	18.56%
	{90970}	{15735}	{10061}	{116766}	{381304}	{59532}	{31884}	{472720}
	\$39.36	\$53.88	\$81.48	\$45.25	\$31.07	\$58.10	\$88.83	\$38.73
Subtotal	13.98%	15.20%	16.75%	14.40%	16.28%	17.49%	18.35%	16.58%
	{627047}	{117701}	{73918}	{818666}	{799155}	{125777}	{73497}	{998429}

This table presents summary statistics of the S&P 500 index option data after applying our filtering to the OptionMetrics Database. The sample ranges from January 4, 1996 to June 28, 2019. The columns Short, Medium and Long refer to the maturity categories. For each moneyness (S/X) and maturity category, the first row depicts the average option price, the second row the average implied volatility and the third row the number of observations (in braces). The average of the daily values of the S&P 500 index and the (annualized) risk-free rate in the sample period were 1443.61 and 2.17%, respectively.

Table 1.2: **Option Price Bounds for S&P 500 Options (1996-2019)**

Category	MDNA Bounds In	MDNA Bounds Upper	MDNA Bounds Lower	MD Bounds In	MD Bounds Upper	MD Bounds Lower	SSD Bounds In	SSD Bounds Upper	SSD Bounds Lower
Panel A: Calls									
Short OTM	27.00%	1.20% (1.47)	71.80% (0.84)	98.37%	1.20% (1.98)	0.43% (0.31)	99.19%	0.60% (2.14)	0.21% (0.28)
Short ATM	80.82%	2.02% (1.19)	17.17% (0.93)	97.98%	2.02% (1.28)	0.00% (0.69)	99.53%	0.47% (1.36)	0.00% (0.67)
Short ITM	94.06%	5.67% (1.06)	0.27% (0.96)	94.25%	5.67% (1.06)	0.07% (0.93)	98.18%	1.75% (1.10)	0.07% (0.93)
Medium OTM	42.82%	0.26% (1.52)	56.92% (0.88)	99.71%	0.26% (2.02)	0.03% (0.31)	99.94%	0.04% (2.26)	0.02% (0.29)
Medium ATM	96.34%	0.45% (1.26)	3.20% (0.91)	99.55%	0.45% (1.27)	0.00% (0.68)	99.98%	0.2% (1.40)	0.00% (0.67)
Medium ITM	97.73%	1.85% (1.11)	0.42% (0.94)	98.14%	1.85% (1.11)	0.01% (0.87)	99.92%	0.30% (1.20)	0.00% (0.87)
Long OTM	74.77%	0.08% (1.61)	25.14% (0.88)	99.91%	0.08% (1.79)	0.00% (0.37)	100%	0.00% (2.07)	0.00% (0.35)
Long ATM	99.84%	0.01% (1.33)	0.15% (0.91)	100%	0.00% (1.33)	0.00% (0.70)	100%	0.00% (1.51)	0.00% (0.68)
Long ITM	99.25%	0.38% (1.18)	0.37% (0.91)	99.60%	0.38% (1.18)	0.01% (0.83)	99.98%	0.00% (1.32)	0.02% (0.82)
All Calls	71.23%	1.88% (1.23)	26.89% (0.92)	98.02%	1.88% (1.46)	0.09% (0.61)	99.42%	0.53% (1.57)	0.05% (0.59)
Panel B: Puts									
Short OTM	97.21%	2.41% (1.85)	0.38% (0.62)	97.53%	2.41% (1.85)	0.05% (0.42)	99.44%	0.49% (2.67)	0.06% (0.42)
Short ATM	83.98%	2.96% (1.21)	13.07% (0.88)	96.17%	2.96% (1.21)	0.88% (0.75)	98.42%	0.80% (1.31)	0.79% (0.74)
Short ITM	28.29%	4.22% (1.02)	67.50% (0.99)	88.85%	4.22% (1.04)	6.93% (0.96)	91.63%	2.48% (1.05)	5.88% (0.95)
Medium OTM	97.74%	1.53% (1.52)	0.73% (0.72)	98.35%	1.53% (1.52)	0.12% (0.50)	99.78%	0.09% (1.95)	0.12% (0.50)
Medium ATM	92.60%	1.40% (1.23)	5.99% (0.90)	98.39%	1.40% (1.23)	0.21% (0.72)	99.56%	0.24% (1.36)	0.20% (0.71)
Medium ITM	51.67%	1.54% (1.07)	46.79% (0.97)	95.99%	1.54% (1.09)	2.46% (0.90)	97.13%	0.52% (1.13)	2.35% (0.90)
Long OTM	97.29%	0.94% (1.47)	1.77% (0.76)	98.67%	0.94% (1.47)	0.40% (0.55)	99.48%	0.11% (1.89)	0.40% (0.55)
Long ATM	92.60%	0.81% (1.30)	6.59% (0.90)	98.69%	0.81% (1.30)	0.50% (0.71)	99.30%	0.20% (1.50)	00.50% (0.71)
Long ITM	68.75%	1.29% (1.14)	29.96% (0.96)	97.36%	1.29% (1.16)	1.35% (0.85)	98.32%	0.34% (1.23)	1.34% (0.85)
All Puts	87.46%	2.48% (1.51)	10.06% (0.76)	96.72%	2.48% (1.47)	0.80% (0.61)	98.63%	0.65% (1.88)	0.72% (0.60)

This table presents empirical results for the minimum dispersion no-arbitrage (MDNA), minimum dispersion (MD) and stochastic dominance (SSD) option price bounds for S&P 500 options. For each category, column In reports the percentage of prices inside the bounds. The column Upper reports the percentage of violations of the upper bounds and, in parenthesis, the average ratio of the upper bound over the observed option price for the prices inside the bounds. Analogously, Lower reports the percentage of violations of the lower bounds and, in parenthesis, the average ratio of the lower bound over the observed option price for the prices inside the bounds. The sample ranges from January 4, 1996 to June 28, 2019.

Table 1.3: **S&P 500 Options Implied γ (1996-2019)**

	Call Options			Put Options		
Moneyness	Short	Medium	Long	Short	Medium	Long
$[0.90, 0.97)$	11.42	8.52	4.51	11.65	7.39	5.27
$[0.97, 1.03)$	1.72	0.29	-0.02	1.33	0.54	0.85
$[1.03, 1.10]$	-1.04	-0.81	-0.54	-0.89	-0.73	-0.41

This table presents the average implied γ of S&P 500 index options for each moneyness (S/X) and maturity category. The sample ranges from January 4, 1996 to June 28, 2019. The columns Short, Medium and Long refer to the maturity categories.

Table 1.4: **Parameters for the SVCJ model**

Measure	μ or r	κ_v	θ_v	σ_v	ρ	λ	μ_s	σ_s	μ_v	ρ_s
$\mathbb{P}$	0.1396	6.5520	0.0135	0.08	-0.4838	1.5120	-0.0263	0.0289	0.0373	0.0
$\mathbb{Q}$	0.05	14.364	0.0061	0.08	-0.4838	1.5120	-0.0539	0.0578	0.2213	0.0

This table presents the annualized parameters used in the simulations of the SVCJ model. The parameters for the physical measure ($\mathbb{P}$) and risk-neutral measure ($\mathbb{Q}$) correspond to the ones estimated in Eraker, Johannes and Polson (2003) and Broadie, Chernov and Johannes (2007), respectively.

1.11 Figures

Figure 1.1: **Absolute Risk Aversion and Prudence of HARA Investors**

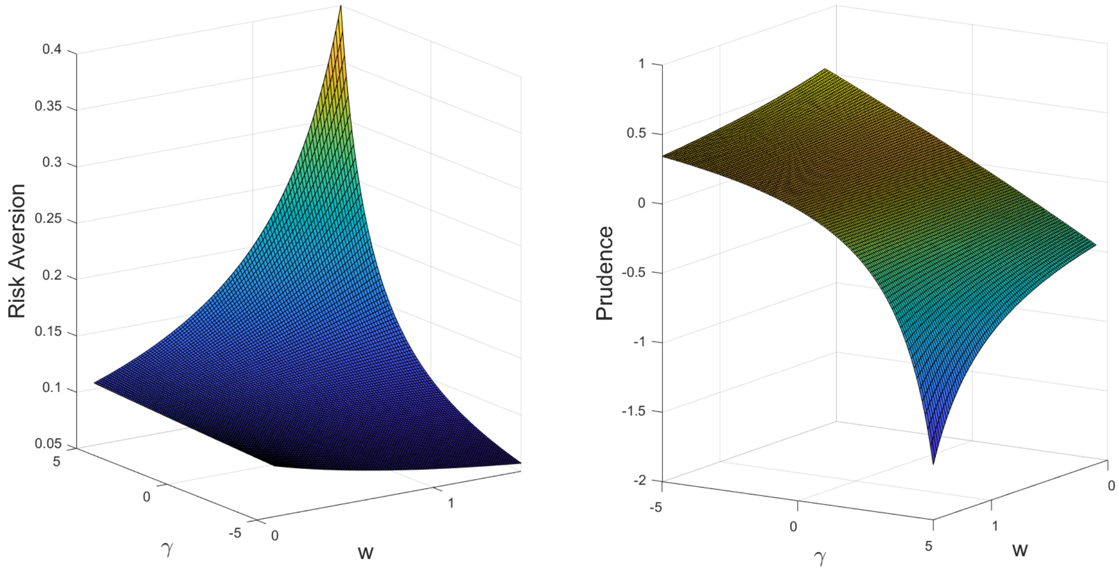

This figure plots the absolute risk aversion and absolute prudence for HARA investors with γ's between -5 and 5 and different levels of wealth W. We set $b = 10$ and $a = 1$ in the utility function.

Figure 1.2: **Minimum Dispersion Risk-Neutral Measures**

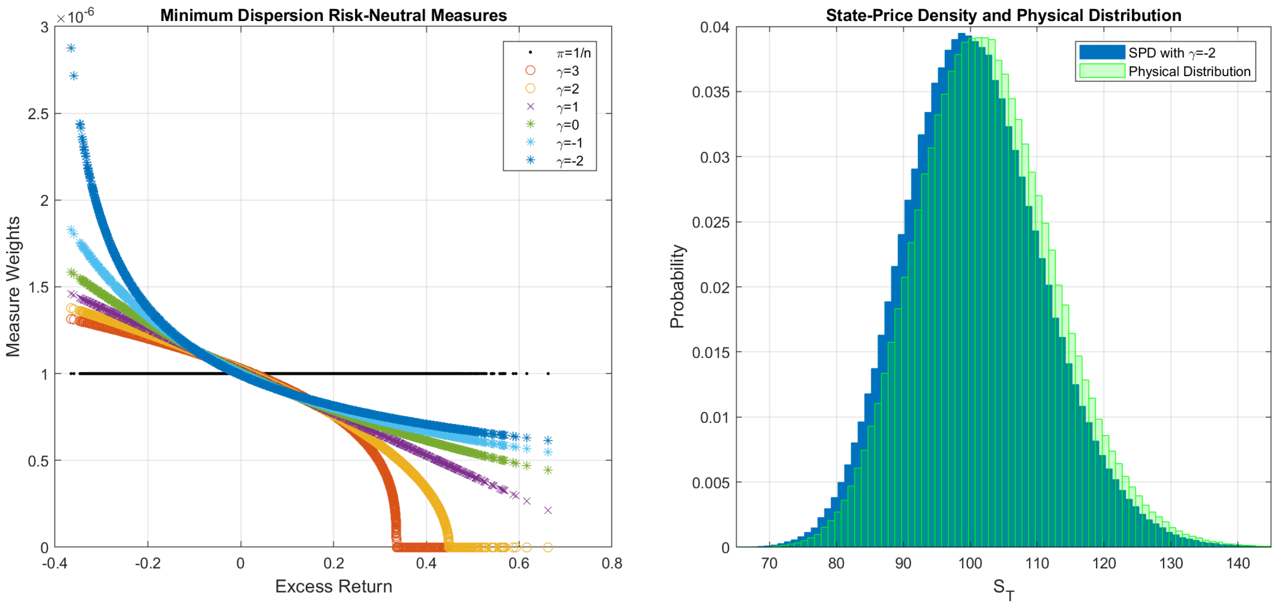

The left panel of this figure plots minimum dispersion risk-neutral measures for different γ's and the empirical measure (π) for one million 3-month returns coming from the Black-Scholes model described in Appendix B 1.8. The right panel plots the state-price density as the histogram of the returns (times a current price of 100) under the risk-neutral probabilities implied by $\gamma = -2$ and the histogram under the empirical measure.

Figure 1.3: **Minimum Dispersion Option Prices, Bounds and Implied γ**

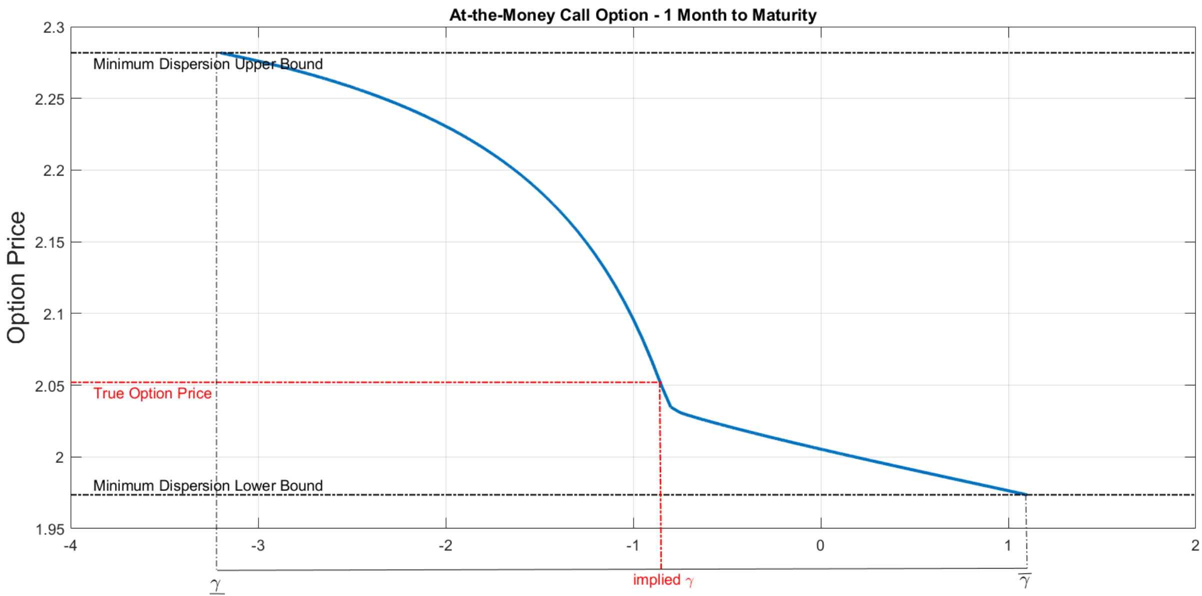

This figure depicts option prices for a European call option with 1 month to maturity and moneyness $S/X = 1$ (ATM) in the SVCJ economy described in Appendix B 1.8. The minimum dispersion admissible set is depicted by the interval $[\underline{\gamma}, \overline{\gamma}]$. The option prices implied by the measures in the set are in blue and the minimum dispersion upper and lower bounds are in black dashed lines. The true option price of the SVCJ model and the implied γ are represented by red dashed lines. The minimum dispersion option prices were calculated from one million 1-month returns sampled from the SVCJ model.

Figure 1.4: **Option Price Bounds - Black-Scholes Economy**

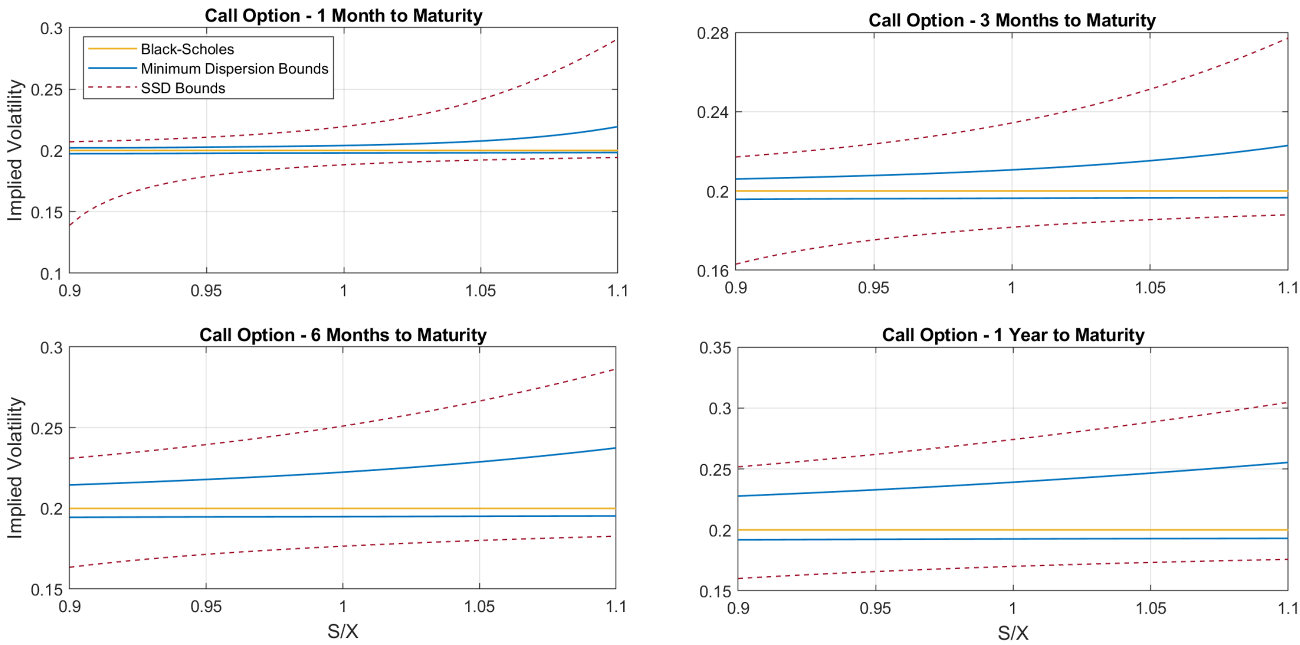

This figure plots the minimum dispersion option price bounds (in blue), stochastic dominance (SSD) bounds (in red dashed lines) and the Black-Scholes theoretical prices (in yellow) converted to Black-Scholes implied volatilities for European call options with different combinations of time to maturity and moneyness (S/X). The bounds were calculated from underlying returns sampled from the physical measure of the economy.

Figure 1.5: **Option Price Bounds - SVCJ Economy**

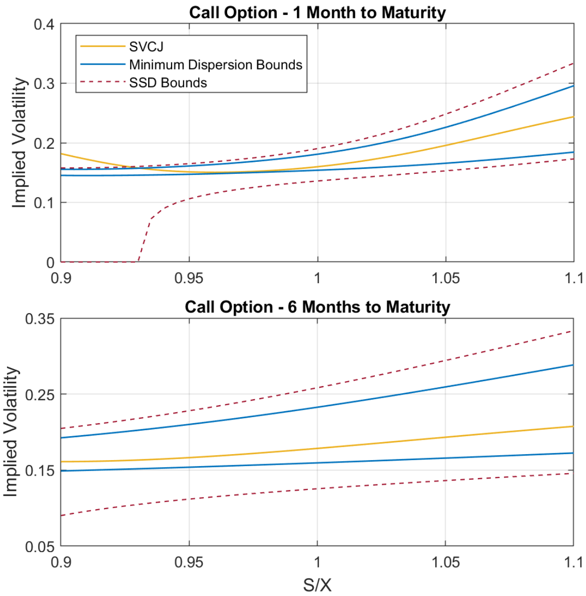

This figure plots the minimum dispersion option price bounds (in blue), stochastic dominance (SSD) bounds (in red dashed lines) and the SVCJ theoretical prices (in yellow) converted to Black-Scholes implied volatilities for European call options with different combinations of time to maturity and moneyness (S/X). The bounds were calculated from underlying returns sampled from the physical measure of the economy.

Figure 1.6: **State-Price Density and Implied γ Surface - Black-Scholes Economy**

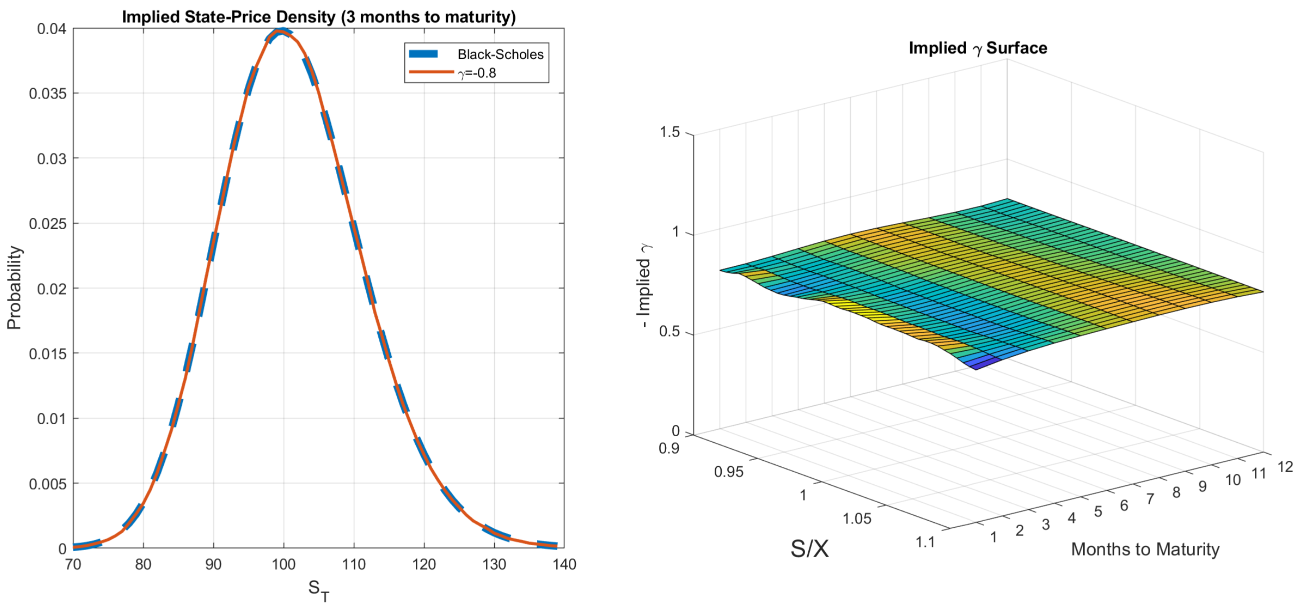

The left panel of this figure plots the SPD implied by the minimum dispersion risk-neutral measure $\gamma^* = -0.8$ and the Black-Scholes model SPD for three months to maturity. The right panel plots the (negative of the) implied γ surface for the Black-Scholes economy. For each combination of moneyness (S/X) and time to maturity we identify the minimum dispersion risk-neutral measure that correctly prices the option.

Figure 1.7: **State-Price Density and Implied γ Surface - SVCJ Economy**

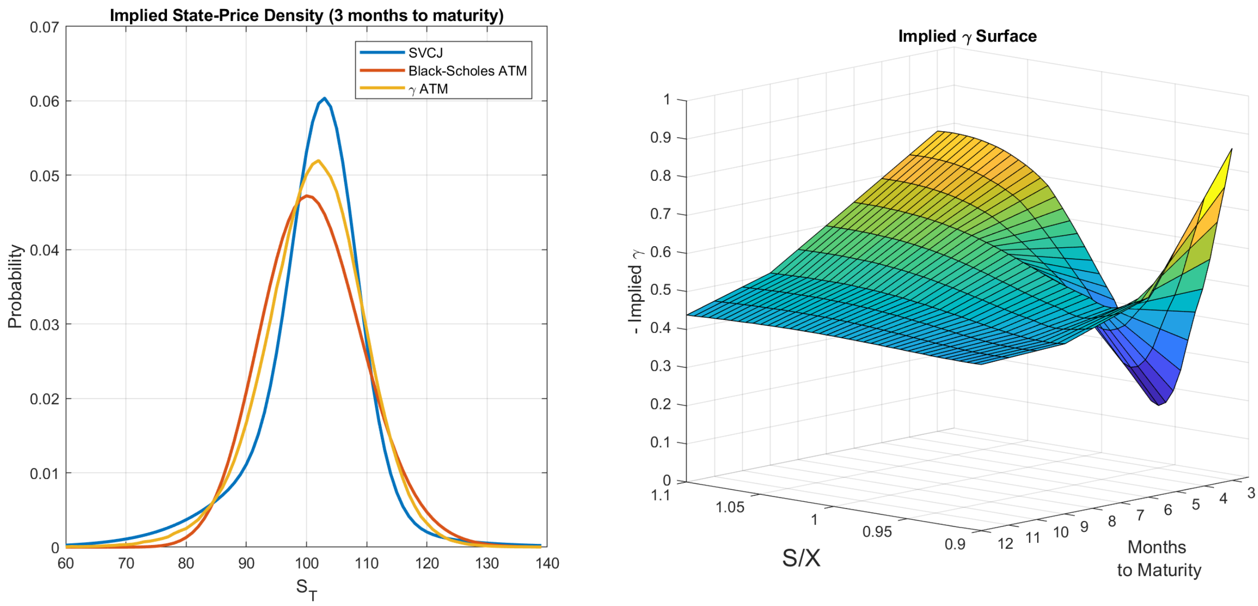

The left panel of this figure plots the SPD implied by the SVCJ model, the SPD implied by the minimum dispersion risk-neutral measure pricing the ATM option and the Black-Scholes model SPD pricing the ATM option for three months to maturity. The right panel plots the (negative of the) implied γ surface for the SVCJ economy. For each combination of moneyness (S/X) and time to maturity we identify the minimum dispersion risk-neutral measure that correctly prices the option.

Figure 1.8: **S&P 500 Options Implied γ Over Time (1996-2019)**

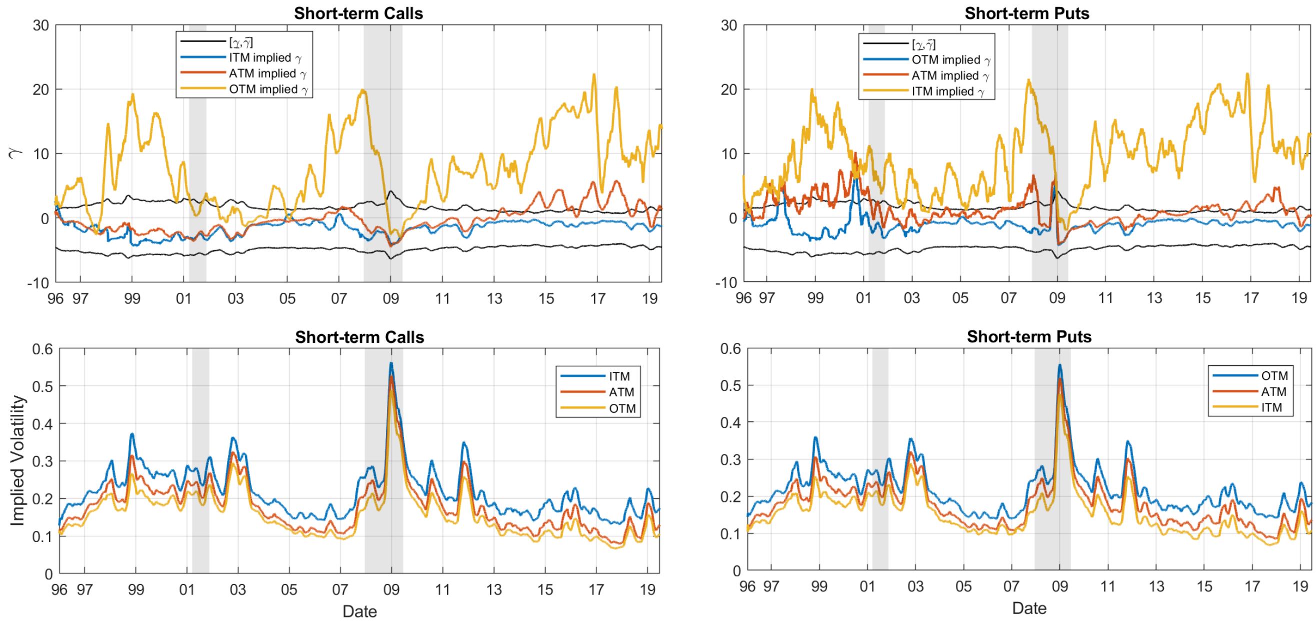

This figure plots, in the upper panels, the 2-month moving averages of the mean implied γ for OTM, ATM and ITM options and the mean $\underline{\gamma}$ and $\overline{\gamma}$ defining the MDNA bounds, for short-term calls and puts. In the lower panels, the corresponding moving averages for mean OTM, ATM and ITM implied volatilities are plotted. Shaded areas depict NBER recession dates. The sample ranges from January 4, 1996 to June 28, 2019.

Figure 1.9: **S&P 500 Option Price Bounds and Implied γ**

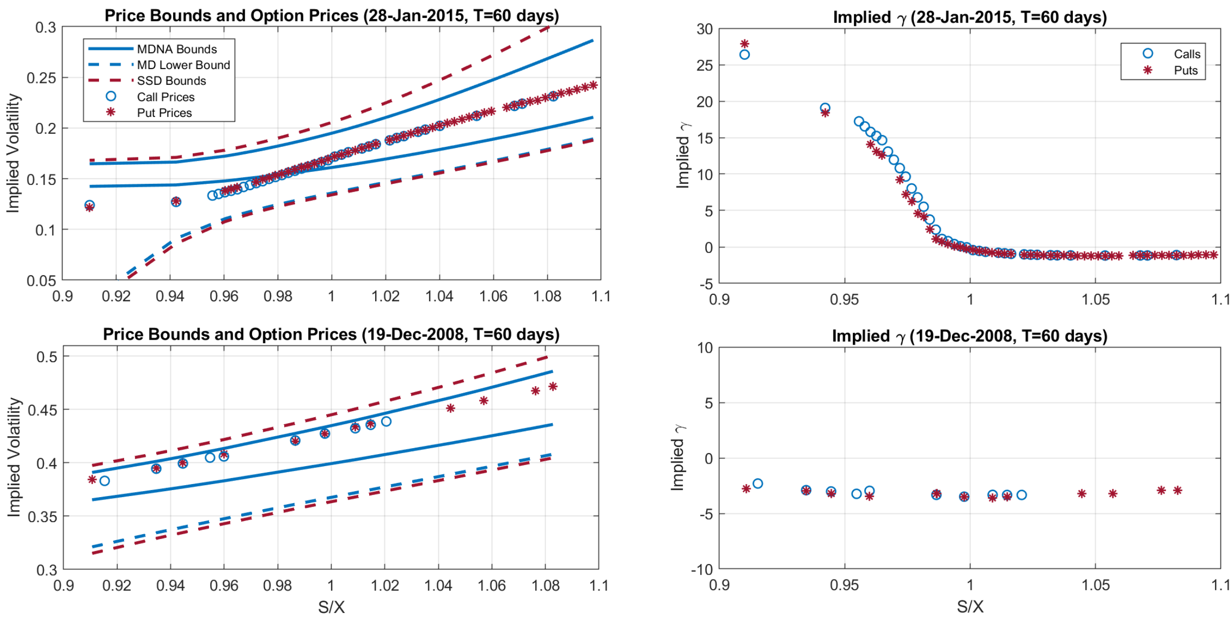

This figure plots the MDNA, MD and SSD bounds and observed S&P 500 option prices converted to implied volatilities in the left panels and the corresponding implied γ in the right panels, for options with 60 days to maturity at two different dates.

Figure 1.10: **S&P 500 Options Implied γ Over Time (1987-1995)**

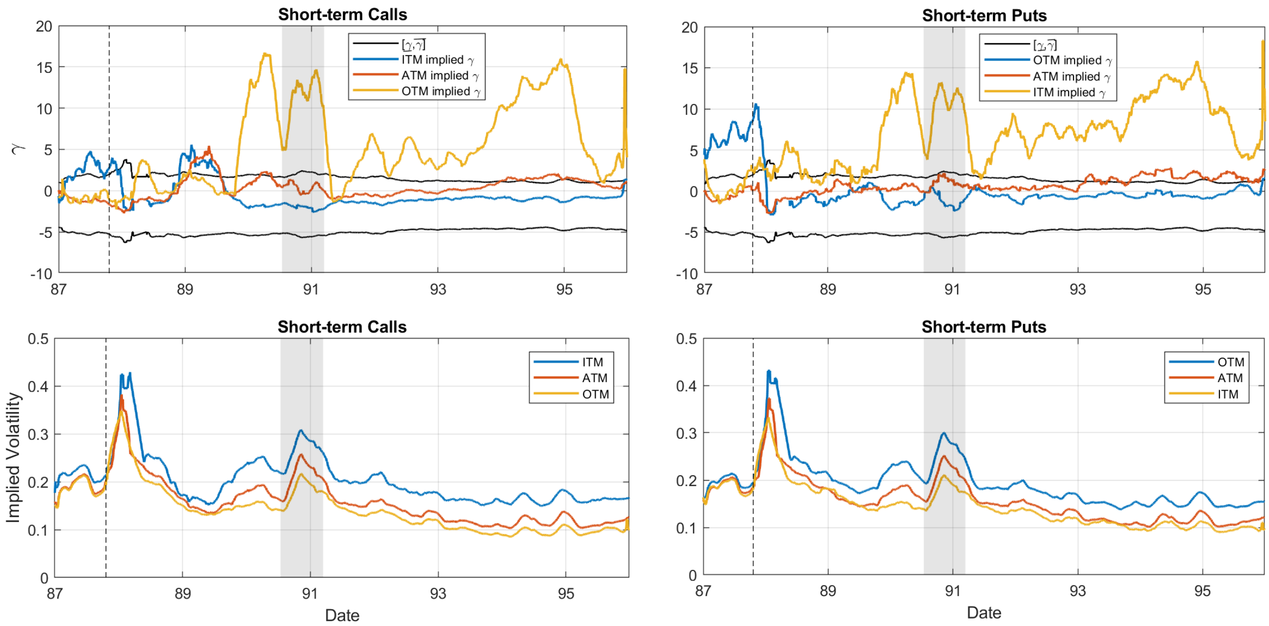

This figure plots, in the upper panels, the 2-month moving averages of the mean implied γ for OTM, ATM and ITM options and the mean $\underline{\gamma}$ and $\overline{\gamma}$ defining the MDNA bounds, for short-term calls and puts. In the lower panels, the corresponding moving averages for mean OTM, ATM and ITM implied volatilities are plotted. Shaded areas depict NBER recession dates and the vertical dashed line corresponds to the October 1987 market crash. The sample ranges from January 2, 1987 to December 29, 1995.

Figure 1.11: **S&P 500 Option Price Bounds and Implied γ Before the 1987 Crash**

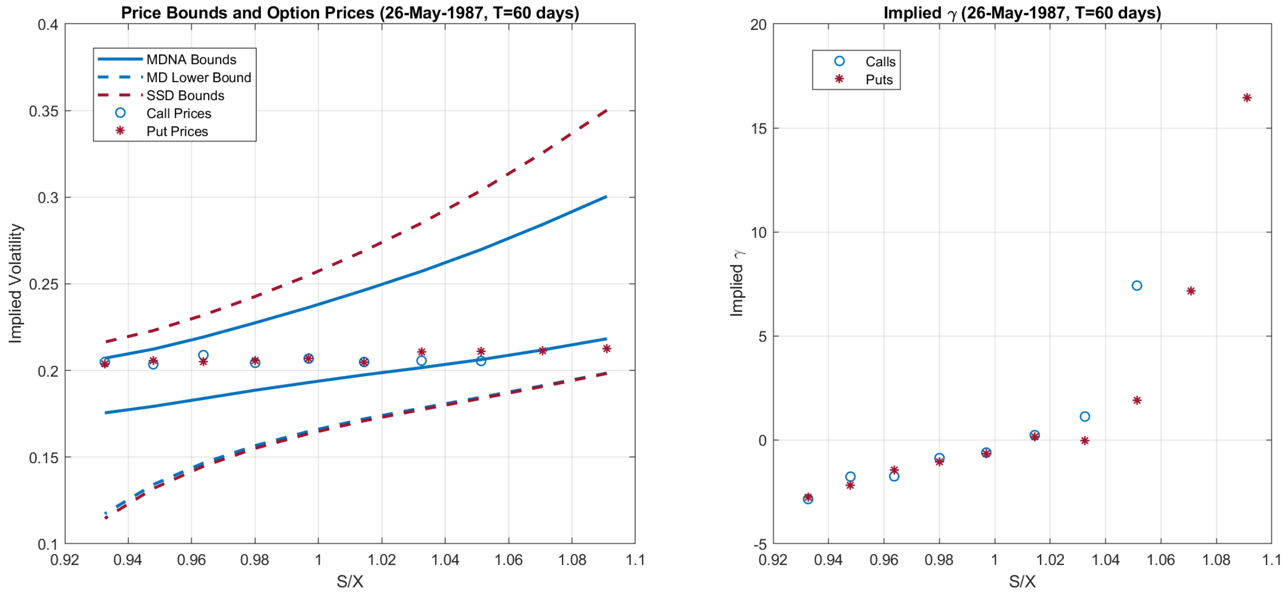

This figure plots the MDNA, MD and SSD bounds and observed S&P 500 option prices converted to implied volatilities in the left panel and the corresponding implied γ in the right panel, for options with 60 days to maturity at May 26, 1987.

Figure 1.12: **Cressie-Read Discrepancies Weights on Higher-Order Moments**

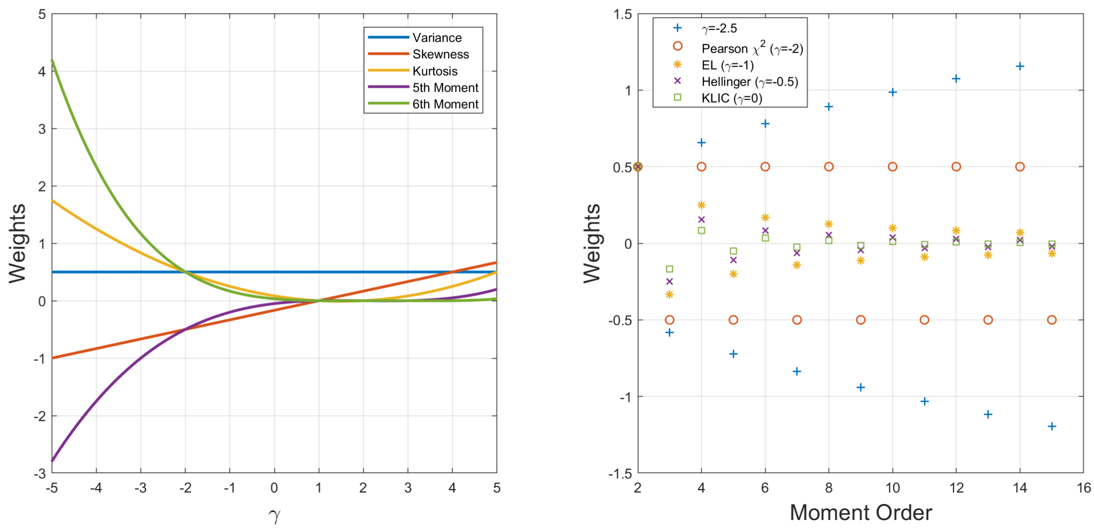

This figure depicts weights given by Cressie-Read discrepancies to higher-order moments of the state-price density. The left panel plots the weights given to the first six moments by discrepancies with γ's between -5 and 5. The right panel plots the weights given to the first 15 moments by specific discrepancies.

1.12 Online Appendix

In Section 1.12.1 of this Online Appendix, we discuss in detail the estimation of minimum dispersion risk-neutral measures. Section 1.12.2 presents theoretical results and proofs that are not shown in the main paper. In Section 1.12.3, we report and analyze additional empirical results on the pricing of S&P 500 options.

1.12.1 Estimation of Minimum Dispersion Risk-Neutral Measures

Consider the environment described in Section 1.2 of the main paper. We start with a general convex function $\phi(.)$ to measure the dispersion between the physical measure $\mathbb{P}$ and the risk-neutral measure $\mathbb{Q}$. A minimum dispersion risk-neutral measure solves the following minimization in the space of admissible measures with $I^{\phi}(\mathbb{Q},\mathbb{P}) \equiv \mathbb{E}\left[\phi(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}})\right] < \infty$:

$$\mathbb{Q}^* = \mathop{arg\,min}_{Q}\, \mathbb{E}\left[\phi\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)\right] \equiv \int \phi\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)\mathrm{d}\mathbb{P}, \ \text{ s.t. } \mathbb{E}^{\mathbb{Q}}\left[\mathbf{R}-\mathbf{R}_f\right] = \mathbf{0}_K, \tag{1.32}$$

where, under the restriction that $\mathbb{Q}$ is absolutely continuous with respect to $\mathbb{P}$, $\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}$ is a Radon-Nikodym derivative. The risk-neutral measure must also be nonnegative and integrate to one. By definition, we have that any admissible measure $\mathbb{Q}$ satisfies $I^{\phi}(\mathbb{Q},\mathbb{P}) \geq I^{\phi}(\mathbb{Q}^*,\mathbb{P})$.

While at first glance the variational problem (1.32) might seem difficult to solve, we follow Kitamura (2006) and Almeida and Garcia (2017) and make use of results in Borwein and Lewis (1991) to show that a duality result comes to rescue:

Theorem 1. *Consider the primal problem:*

$$\min_{Q}\, \mathbb{E}\left[\phi(\tfrac{d\mathbb{Q}}{d\mathbb{P}})\right], \ \ s.t. \ \mathbb{E}^{\mathbb{Q}}\left[\boldsymbol{R}-\boldsymbol{R}_f\right] = \boldsymbol{0}_K, \ \ \mathbb{E}(\tfrac{d\mathbb{Q}}{d\mathbb{P}}) = 1, \ \ \tfrac{d\mathbb{Q}}{d\mathbb{P}} \geq 0, \tag{1.33}$$

and the dual problem:

$$\max_{\alpha\in R,\, \lambda\in R^K}\ \alpha - \mathbb{E}\left[\, \phi^{*,+}(\alpha+\lambda'(\boldsymbol{R}-\boldsymbol{R}_f)) + \delta([\alpha\ \lambda]\,|\,\Lambda(\boldsymbol{R}))\,\right], \tag{1.34}$$

where $\Lambda(\boldsymbol{R}) = \{\alpha\in\mathbb{R},\, \lambda\in\mathbb{R}^K : (\alpha+\lambda'(\boldsymbol{R}-\boldsymbol{R}_f)) \in dom\ \phi^{*,+}\}$,[48] $\delta(.\,|\,C)$ *is such that*

[48]The domain of $\phi^{*,+}(z)$ is defined as the values of z for which the function is finite ($\phi^{*,+}(z) < \infty$).

$\delta(x\,|\,C) = 0$ if $x \in C$ and ∞ otherwise, and $\phi^{*,+}$ denotes the convex conjugate of ϕ:

$$\phi^{*,+}(z) = \underset{w\in[0,\infty)\cap domain\ \phi}{sup} zw - \phi(w). \tag{1.35}$$

Absence of arbitrage implies that the values of the primal and the dual problems coincide (with dual attainment). A sufficient condition allowing the unique minimum dispersion risk-neutral measure to be obtained from the solution of the dual optimization problem is that either $d = lim_{x\to\infty} \frac{\phi(x)}{x} = \infty$ or ($d < \infty$ and $c = lim_{x\to\infty} ((d - \phi'(x))x) > 0$). In such cases, the implied state-price density is obtained by:

$$\frac{d\mathbb{Q}^*}{d\mathbb{P}} = \phi^{*,+}(z)z_{z=\left(\alpha^*+\lambda^{*\prime}(\boldsymbol{R}-\boldsymbol{R}_f)\right)}, \tag{1.36}$$

where α^* and λ^* are the optimal Lagrange multipliers solving the dual problem (1.34).

Proof. Let $\overline{m} \equiv m/\mathbb{E}(m)$ be the normalized version of a given SDF m. When m is nonnegative, $\overline{m}$ will generate a state-price density. In Theorem 2.4 of Borwein and Lewis (1991), let X be the space of normalized SDFs $\overline{m}$ such that $\mathbb{E}[\overline{m}(\mathbf{R} - \mathbf{R}_f)] = 0$ and $\mathbb{E}[\phi(\overline{m})] < \infty$, $f(\overline{m}) = \mathbb{E}[\phi(\overline{m})]$, $C = X^+$ be the space of normalized nonnegative SDFs $\overline{m}$ such that $\mathbb{E}[\overline{m}(\mathbf{R} - \mathbf{R}_f)] = 0$ and $\mathbb{E}[\phi(\overline{m})] < \infty$, $A\overline{m} = \mathbb{E}(\overline{m}[(\mathbf{R} - \mathbf{R}_f)'\ 1]')$, $b = [\mathbf{0}\ 1]'$ and $P = 0$. Theorem 2.5 at page 327 in Borwein and Lewis (1991) allows us to conjugate $\phi(.)$ within the expectation to obtain $g^* = \mathbb{E}(\phi^{*,+})$. In addition, we obtain $A'\tilde{\lambda} = \tilde{\lambda}'[(\mathbf{R} - \mathbf{R}_f)'\ 1]'$ and $P^+ = \mathbb{R}^K$. No-arbitrage guarantees the existence of at least one feasible point (a strictly positive normalized SDF such that $\mathbb{E}[\overline{m}(\mathbf{R} - \mathbf{R}_f)] = 0$) in the quasi-relative interior of X^+. Therefore, rewriting $\tilde{\lambda} = [\alpha\ \lambda]$, the primal problem in Theorem 1 has a solution that coincides with that of the following dual problem:

$$\underset{\tilde{\lambda}\in R^{K+1}}{max}\ \alpha - \mathbb{E}\left[\phi^{*,+}(\alpha + \lambda'(\mathbf{R} - \mathbf{R}_f)) + \delta(\tilde{\lambda}\,|\,\Lambda(\mathbf{R}))\right] \tag{1.37}$$

where $\Lambda(\mathbf{R}) = \{\tilde{\lambda} \in \mathbb{R}^{K+1} : (\alpha + \lambda'(\mathbf{R} - \mathbf{R}_f)) \in \text{dom}\ \phi^{*,+}\}$. Furthermore, if either $d = lim_{x\to\infty} \frac{\phi(x)}{x} = \infty$ or ($d < \infty$ and $c = lim_{x\to\infty} ((d - \phi'(x))x) > 0$), Theorem 5.5 at page 335 of Borwein and Lewis (1991) guarantees that the unique primal optimal solution is obtained by differentiating the convex conjugate $\phi^{*,+}(z)$ at the optimal (Lagrange multipliers) dual solution: $\overline{m}^* = \mathrm{d}\mathbb{Q}^*/\mathrm{d}\mathbb{P} = \phi^{*,+}(z)z_{z=\left(\alpha^*+\lambda^{*\prime}(\mathbf{R}-\mathbf{R}_f)\right)}$. □

The dual problem is a much simpler finite dimensional convex maximization problem. The absence of arbitrage implies the existence of an interior point in the space of admissible state-price densities. This is a fundamental condition to guarantee that the solutions of the primal and dual problems coincide. The vector λ and the scalar α are Lagrange multipliers coming from the primal problem constraints given by the Euler equations for the basis assets and the restriction $\mathbb{E}(\mathrm{d}\mathbb{Q}/\mathrm{d}\mathbb{P}) = 1$, respectively. The primal problem nonnegativity constraint restricts the convex conjugate to be calculated on the nonnegative real line, while the delta function $\delta(.\,|\,\Lambda(\mathbf{R}))$ restricts, for each vector of returns $\mathbf{R}$ in the probability space, the optimization problem to a subset $\Lambda(\mathbf{R})$ of $\mathbb{R}^K$ where the convex conjugate assumes finite values.

Theorem 1 provides a general result for obtaining a minimum discrepancy risk-neutral measure according to a convex function ϕ, which can be seen as a generalized minimum contrast estimation procedure (Bickel et al., 1993). This procedure falls into the category of methods that treat the data distribution nonparametrically in statistical estimation by comparing the empirical distribution of the data with the family of distributions implied by a statistical model. In our case, the statistical model is defined to be the set of all probability measures that satisfy the moment restrictions characterizing risk neutrality for the basis assets.

We focus on risk-neutral measures minimizing the Cressie and Read (1984) family of discrepancies, defined as:

$$\phi_\gamma\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right) = \frac{\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)^{\gamma+1} - 1}{\gamma(\gamma+1)}, \quad \gamma \in \mathbb{R}. \tag{1.38}$$

Kitamura (1996), Baggerly (1998) and Newey and Smith (2004) suggest using this comprehensive divergence family, that includes as particular cases several well-known measures of dispersion. For instance, the Euclidean divergence ($\gamma = 1$), the KLIC ($\gamma \to 0$), the Hellinger divergence ($\gamma = -1/2$), the empirical likelihood ($\gamma \to -1$) and Pearson's Chi-Square ($\gamma = -2$). Furthermore, Newey and Smith (2004) show that the minimum discrepancy Cressie-Read estimators are equivalent to the class of generalized empirical likelihood (GEL) estimators (Smith, 1997). That is, the GEL objective function is given by the dual problem of the minimum discrepancy problem, where γ indexes the particular estimator in this class. For instance, $\gamma \to -1$ yields the empirical likelihood (Owen, 1988), $\gamma \to 0$ the exponential tilting (Kitamura and Stutzer, 1997) and $\gamma = 1$ the continuous updating estimator (Hansen, Heaton and Yaron, 1996). These estimators are robust againts distributional assumptions, possess desirable properties analogous to those of parametric likelihood procedures and have been used to improve on the small sample properties of GMM estimators.[49]

By considering the Cressie-Read family, Theorem 1 produces the following corollary:

Corollary 2. *Let the discrepancy function ϕ in the minimization problem (1.33) belong to the family defined in (1.38), and assume absence of arbitrage. Then, letting $\Lambda_\gamma(\boldsymbol{R}) = \{\lambda \in \mathbb{R}^K \,:\, (1 + \gamma\lambda'(\boldsymbol{R} - \boldsymbol{R}_f)) > 0\}$:*

[49] See Kitamura (2006) for a review.

i) if $\gamma > 0$*, (1.34) specializes to:*

$$\alpha_\gamma^* = \frac{1}{\gamma},\ \lambda_\gamma^* = \underset{\lambda \in R^K}{arg\ max}\ \ -\frac{1}{\gamma+1}\mathbb{E}\left[(1+\gamma\lambda'(\boldsymbol{R}-\boldsymbol{R}_f))^{\left(\frac{\gamma+1}{\gamma}\right)} I_{\Lambda_\gamma(\boldsymbol{R})}(\lambda)\right]. \tag{1.39}$$

ii) if $\gamma < 0$*, it specializes to:*

$$\alpha_\gamma^* = \frac{1}{\gamma},\ \lambda_\gamma^* = \underset{\lambda \in R^K}{arg\ max}\ \ -\frac{1}{\gamma+1}\mathbb{E}\left[(1+\gamma\lambda'(\boldsymbol{R}-\boldsymbol{R}_f))^{\left(\frac{\gamma+1}{\gamma}\right)} - \delta(\lambda \,|\, \Lambda_\gamma(\boldsymbol{R}))\right]. \tag{1.40}$$

iii) if $\gamma = 0$*, the maximization is unconstrained:*

$$\alpha_0^* = 1,\ \lambda_0^* = \underset{\lambda \in R^K}{arg\ max}\ \ -\mathbb{E}\left[e^{\lambda'\left(\boldsymbol{R}-\boldsymbol{R}_f\right)}\right], \tag{1.41}$$

where $I_A(x) = 1$ *if* $x \in A$*, and* 0 *otherwise.*

Proof. We need to obtain the convex conjugate ϕ_γ^* of ϕ_γ belonging to the Cressie-Read family as in (1.38) to substitute in the dual problem (1.34). Given equation (1.35) defining the convex conjugate, we define an auxiliary function $h_\gamma^z(w) = zw - \frac{w^{\gamma+1}-a^{\gamma+1}}{\gamma(\gamma+1)}$, whose domain is $dom(h_\gamma^z) = [0,\infty) \cap \mathrm{dom}(\phi_\gamma)$. Note that for $\gamma > -1$ and $\gamma \neq 0$, $\mathrm{dom}(h_\gamma^z) = [0,\infty)$, and for $\gamma \leq -1$ or $\gamma = 0$, $\mathrm{dom}(h_\gamma^z) = (0,\infty)$. In order to obtain the supremum in $\phi_\gamma^{*,+}(z) = sup_{w\in\mathrm{dom}(h_\gamma^z)} h_\gamma^z(w)$, we differentiate $h_\gamma^z(w)$ with respect to w, leading to: $\frac{dh_\gamma^z(\tilde{w})}{dw} = z - \frac{\tilde{w}^\gamma}{\gamma}$. Now we split the analysis in three cases: $\gamma > 0$, $\gamma < 0$, and $\gamma = 0$.

i) $\gamma > 0$.

In this case, $dom(h_\gamma^z) = [0,\infty)$. If $z \leq 0$, h_γ^z is a decreasing function of w and achieves its maximum at $\tilde{w} = 0$. If $z > 0$, $\tilde{w} = (\gamma z)^{\frac{1}{\gamma}}$ will be the unique critical point where the function achieves its maximum. By combining these two solutions, we note that $dom(\phi_\gamma^{*,+}) = \mathbb{R}$ (which implies that $\Lambda_\gamma = \mathbb{R}^{K+1}$), and, for an arbitrary z, we obtain $\tilde{w} = (\gamma z)^{\frac{1}{\gamma}} \mathbb{1}_{\{\gamma z > 0\}}$. Substituting $\tilde{w}$ in $\phi_\gamma^{*,+}(z) = h_\gamma^z(\tilde{w})$, the convex conjugate becomes:

$$\phi_\gamma^{*,+}(z) = \frac{(\gamma z)^{\frac{\gamma+1}{\gamma}}}{\gamma+1}\mathbb{1}_{\{\gamma z>0\}} + \frac{1}{\gamma(\gamma+1)}, \tag{1.42}$$

and the optimization problem becomes:

$$\tilde{\lambda}_\gamma^* = \underset{\alpha\in R,\, \lambda \in R^K}{arg\ max}\ \alpha - \mathbb{E}\left[\frac{\left(\gamma\left(\alpha+\lambda'\left(\mathbf{R}-\mathbf{R}_f\right)\right)\right)^{\frac{\gamma+1}{\gamma}}}{\gamma+1}\mathbb{1}_{\left\{\gamma\left(\alpha+\lambda'\left(\mathbf{R}-\mathbf{R}_f\right)\right)>0\right\}}\right] - \frac{1}{\gamma(\gamma+1)}. \tag{1.43}$$

First, note that we can discard the constant $\frac{1}{\gamma(\gamma+1)}$ from the maximization problem above. Now, as Kitamura (2006, page 12) notes, the convex conjugate ϕ_γ^* is homogeneous, so we can concentrate out α and re-define λ as $\overline{\lambda} = \lambda/\alpha$ to obtain:

$$\alpha - \frac{(\gamma\alpha)^{\frac{\gamma+1}{\gamma}}}{\gamma+1}\mathbb{E}\left[\left(1+\overline{\lambda}'\left(\mathbf{R}-\mathbf{R}_f\right)\right)^{\frac{\gamma+1}{\gamma}} \mathbb{1}_{\{\gamma\alpha\left(1+\overline{\lambda}'\left(\mathbf{R}-\mathbf{R}_f\right)\right)>0\}}\right]. \tag{1.44}$$

Ignoring for now the indicator function inside the expectation, let $\Gamma(\alpha) = \alpha - \frac{(\gamma\alpha)^{\frac{\gamma+1}{\gamma}}}{\gamma+1}$, where the optimal concentrated α is obtained by maximizing Γ. From its first-order condition, we get: $\frac{d\Gamma(\alpha)}{d\alpha} = 0 \Rightarrow \alpha^*_\gamma = \frac{1}{\gamma}$. Since $\gamma\alpha^*_\gamma = 1 > 0$, it does not affect the indicator function. Substituting α^*_γ in $\overline{\lambda}$ and (1.44), we have:

$$\lambda^*_\gamma = \underset{\lambda \in R^K}{arg\ max}\ \frac{1}{\gamma} - \frac{1}{\gamma+1}\mathbb{E}\left[\left(1+\gamma\lambda'\left(\mathbf{R}-\mathbf{R}_f\right)\right)^{\frac{\gamma+1}{\gamma}} \mathbb{1}_{\{1+\gamma\lambda'\left(\mathbf{R}-\mathbf{R}_f\right)>0\}}\right]. \tag{1.45}$$

Since the first term does not affect the maximization, we have the desired result.

ii) $\gamma < 0$.

In this case, $dom(h^z_\gamma) = [0,\infty)$ if $-1 < \gamma < 0$ and $dom(h^z_\gamma) = (0,\infty)$ if $\gamma \leq -1$. If $z \geq 0$, h^z_γ is an increasing function of w and achieves its maximum at $\tilde{w} = \infty$. If $z < 0$, $\tilde{w} = (\gamma z)^{\frac{1}{\gamma}}$ will be the unique critical point where the function achieves its maximum. The fact that $\phi^{*,+}_\gamma$ is ∞ for $z \geq 0$ and is finite otherwise directly implies that $dom(\phi^{*,+}_\gamma) = (-\infty, 0)$. Combining these two solutions, the convex conjugate becomes:

$$\phi^{*,+}_\gamma(z) = \frac{(\gamma z)^{\frac{\gamma+1}{\gamma}}}{\gamma+1} + \delta(z|\{\tilde{z} \in \mathbb{R} : \gamma\tilde{z} > 0\}) + \frac{1}{\gamma(\gamma+1)}, \tag{1.46}$$

and the optimization problem becomes:

$$\tilde{\lambda}^*_\gamma = \underset{\alpha \in R,\, \lambda \in R^K}{arg\ max}\ \alpha - \mathbb{E}\left[\frac{\left(\gamma\left(\alpha+\lambda'\left(\mathbf{R}-\mathbf{R}_f\right)\right)\right)^{\frac{\gamma+1}{\gamma}}}{\gamma+1} + \delta([\alpha\ \lambda])|\Lambda_\gamma(\mathbf{R}))\right] - \frac{1}{\gamma(\gamma+1)}, \tag{1.47}$$

where $\Lambda_\gamma(\mathbf{R}) = \{\alpha \in \mathbb{R}, \lambda \in \mathbb{R}^K : \gamma\left(\alpha + \lambda'\left(\mathbf{R}-\mathbf{R}_f\right)\right) > 0\}$. Following the same procedure as in the previous case to concentrate out α, we obtain the desired result, with $\Lambda_\gamma(\mathbf{R})$ simplifying to $\Lambda_\gamma(\mathbf{R}) = \{\lambda \in \mathbb{R}^K : (1 + \gamma\lambda'\left(\mathbf{R}-\mathbf{R}_f\right)) > 0\}$.

iii) $\gamma = 0$.

The limit $lim_{\gamma\to 0}\mathbb{E}(\phi_\gamma(\overline{m})) = \mathbb{E}(m\log(\overline{m}))$ coincides with the KLIC (Stutzer, 1995, page 375). Therefore, we need to obtain the convex conjugate $\phi^{*,+}_0$ of $\phi_0(\overline{m}) = m\log(\overline{m})$, whose domain is $\mathrm{dom}(\phi_0) = (0,\infty)$. Note that the corresponding auxiliary function is $h^z_0(w) = zw - w\log(w)$, whose first derivative is $\frac{dh^z_0(\tilde{w})}{dw} = z - 1 - \log(\tilde{w})$. Since $dom(h^z_0) = (0,\infty)$, and in this range $\log(w)$ covers the whole real line, for any value of z the only critical point will be $\tilde{w} = e^{z-1}$, implying that $dom(\phi^{*,+}_0) = \mathbb{R}$ (which in turn implies that $\Lambda_\gamma = \mathbb{R}^{K+1}$). Substituting $\tilde{w}$ in $h_0^z(\tilde{w})$, the convex conjugate becomes:

$$\phi_0^{*,+}(z) = e^{z-1}, \tag{1.48}$$

and the optimization problem becomes:

$$\tilde{\lambda}_0^* = \underset{\alpha\in R,\, \lambda\in R^K}{arg\ max}\ \alpha - \mathbb{E}\left[e^{(\alpha-1)+\lambda'(\mathbf{R}-\mathbf{R}_f)}\right]. \tag{1.49}$$

To concentrate α out of (1.49), we define $\Gamma(\alpha) = \alpha - e^{\alpha-1}$ and obtain its first-order condition, to see that it is maximized at $\alpha_0^* = 1$. Substituting α_0^* in (1.49) we obtain the desired result. □

In the next corollary, we use Theorem 1 to identify the implied minimum dispersion state-price densities and to verify if the different discrepancies of the Cressie-Read family satisfy the regularity sufficient conditions that allow us to obtain the risk-neutral measures from the first derivative of the convex conjugate $\phi^{*,+}(.)$.

Corollary 3. *Let the discrepancy function ϕ in the minimization problem (1.33) belong to the family defined in (1.38). For any $\gamma \geq -1$, at least one of the sufficient regularity conditions stated in Theorem 1 is satisfied by ϕ_γ and the corresponding minimum dispersion implied state-price density will be given by:*

$$\frac{d\mathbb{Q}^*}{d\mathbb{P}}(\gamma, \boldsymbol{R}) = \left(1 + \gamma\lambda_\gamma^{*\prime}(\boldsymbol{R} - \boldsymbol{R}_f)\right)^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\boldsymbol{R})}(\lambda_\gamma^*),\ \gamma > 0, \tag{1.50}$$

$$\frac{d\mathbb{Q}^*}{d\mathbb{P}}(\gamma, \boldsymbol{R}) = \left(1 + \gamma\lambda_\gamma^{*\prime}(\boldsymbol{R} - \boldsymbol{R}_f)\right)^{\frac{1}{\gamma}},\ -1 \leq \gamma < 0, \tag{1.51}$$

$$\frac{d\mathbb{Q}^*}{d\mathbb{P}}(0, \boldsymbol{R}) = e^{\lambda_0^{*\prime}(\boldsymbol{R}-\boldsymbol{R}_f)},\ \gamma = 0. \tag{1.52}$$

where for $\gamma > 0$, λ_γ^* solves (1.39), for $-1 \leq \gamma < 0$, λ_γ^* solves (1.40) and λ_0^* solves (1.41).

For $\gamma < -1$, none of the stated regularity conditions in Theorem 1 are satisfied. In this case, an alternative sufficient condition for the minimum dispersion implied state-price density to be given by the expression in (1.51), with λ_γ^* solving (1.40), is that:

$$inf_{\omega\in\Omega}\ \left(1 + \gamma\lambda_\gamma^{*\prime}(\boldsymbol{R}(\omega) - \boldsymbol{R}_f)\right) > 0. \tag{1.53}$$

If the sample space has a finite number of states, the infimum becomes a minimum, and this condition is equivalent to $\forall\omega \in \Omega : \left(1 + \gamma\lambda_\gamma^{*\prime}(\boldsymbol{R}(\omega) - \boldsymbol{R}_f)\right) > 0$.

Proof. *i)* $\gamma > 0$.

It is easy to see that $d = lim_{x\to\infty}\frac{\phi_\gamma(x)}{x} = lim_{x\to\infty}\frac{x^\gamma - \frac{1}{x}}{\gamma(\gamma+1)} = \infty$, implying that the first sufficient condition stated at Theorem 1 is satisfied.

ii) $-1 < \gamma < 0$.

Since, in this case, $d = lim_{x\to\infty} \frac{x^{\gamma} - \frac{1}{x}}{\gamma(\gamma+1)} = 0$, we proceed to verify if $c = lim_{x\to\infty} ((d - \phi'(x))x) > 0$. As $\phi'(x) = \frac{x^{\gamma}}{\gamma}$, we obtain $c = lim_{x\to\infty} - \frac{x^{\gamma+1}}{\gamma} = \infty > 0$, implying that the second sufficient condition at Theorem 1 is satisfied.

iii) $\gamma = -1$.

Since $\phi_{-1}(x) = -\log(x)$, we have ${\phi_{-1}}'(x) = -\frac{1}{x}$ and $c = lim_{x\to\infty} -(-\frac{1}{x})x = 1 > 0$, implying that the second sufficient condition at Theorem 1 is satisfied.

iv) $\gamma < -1$.

In this case, $d = 0$ and $c = lim_{x\to\infty} - \frac{x^{\gamma+1}}{\gamma} = 0$, implying that none of the two sufficient conditions appearing in Theorem 1 are satisfied. Therefore, we invoke Theorem 4.8 at page 334 of Borwein and Lewis (1991), which adapted to our problem states that if the optimal Lagrange multipliers in the dual problem satisfy:[50]

$$inf_{\omega\in\Omega} \left(1 + \gamma\lambda_{\gamma}^{*\prime} \left(\mathbf{R}(\omega) - \mathbf{R}_f\right)\right) > 0, \tag{1.54}$$

the unique primal optimal solution can be obtained by differentiating the convex conjugate. If the sample space has a finite number of states, the infimum becomes a minimum, and (1.54) is equivalent to $\forall\omega \in \Omega : \left(1 + \gamma\lambda_{\gamma}^{*\prime} \left(\mathbf{R}(\omega) - \mathbf{R}_f\right)\right) > 0$.

Now, we show that, whenever (1.36) is valid, differentiating the convex conjugate in (1.42), (1.46) and (1.48) with respect to z gives:

$$\phi^{*,+}(z)z = (\gamma z)^{\frac{1}{\gamma}} \mathbb{1}_{\{\gamma z \geq 0\}},\ \gamma > 0, \tag{1.55}$$

$$\phi^{*,+}(z)z = (\gamma z)^{\frac{1}{\gamma}},\ \gamma < 0, \tag{1.56}$$

$$\phi^{*,+}(z)z = e^{z-1},\ \gamma = 0. \tag{1.57}$$

Substituting $z = \left(\alpha_{\gamma}^{*} + \lambda_{\gamma}^{*\prime} \left(\mathbf{R} - \mathbf{R}_f\right)\right)$, with $\alpha_{\gamma}^{*} = \frac{1}{\gamma}$ for $\gamma \neq 0$ and $\alpha_{0}^{*} = 1$ for $\gamma = 0$, in (1.55), (1.56) and (1.57), we obtain expressions (1.50), (1.51) and (1.52). $\square$

The results so far focused on the population. In order to estimate minimum dispersion risk-neutral measures from data on basis assets returns, we consider the sample version of problem (1.33). In this case, the sample space Ω is finite and discrete, with states of nature $k = \{1, ..., n\}$, where $n > K$. Let $\{\mathbf{R}_k\}_{k=1}^{n}$ be the observed gross returns of the K basis assets, where each $\mathbf{R}_k$ is independent and identically distributed according to $\mathbb{P}$. The unknown physical measure $\mathbb{P}$ can be replaced by the empirical measure $\mathbb{P}_n$

[50] The original condition appearing at Theorem 4.8 of Borwein and Lewis would be replicated in our problem by $sup_{\omega\in\Omega}\ \alpha_{\gamma}^{*} + \lambda_{\gamma}^{*\prime} \left(\mathbf{R}(\omega) - \mathbf{R}_f\right) < 0$. We transform it by multiplying by γ and substituting the concentrated $\alpha_{\gamma}^{*} = \frac{1}{\gamma}$.

that gives weights $\pi_k = 1/n$ to the realization of each state of nature.[51] This allows us to exchange the expectation $\mathbb{E}$ with its sample counterpart $\frac{1}{n}\sum_{k=1}^{n} \equiv \sum_{k=1}^{n}\pi_k$. In the following corollary (Corollary 1 in the main paper), we summarize the results for the sample version of the problem of finding a minimum dispersion Cressie-Read risk-neutral measure:

Corollary 4. *Consider the primal problem:*

$$\begin{aligned} &\underset{\{\pi_1^Q,\dots,\pi_n^Q\}}{min} \ \sum_{k=1}^{n}\pi_k \frac{\left(\pi_k^{\mathbb{Q}}/\pi_k\right)^{\gamma+1}-1}{\gamma(\gamma+1)}, \\ &s.t.\ \sum_{k=1}^{n}\pi_k^{\mathbb{Q}}\left(\boldsymbol{R}_k-\boldsymbol{R}_f\right)=0,\ \ \sum_{k=1}^{n}\pi_k^{\mathbb{Q}}=1,\ \pi_k^{\mathbb{Q}}\geq 0\ \ \forall\, k. \end{aligned} \tag{1.58}$$

Absence of arbitrage in the observed sample implies that the value of the primal problem coincides (with dual attainment) with the value of the dual problem below:

i) if $\gamma > 0$:

$$\lambda_\gamma^* = \underset{\lambda\in R^K}{arg\ max} -\frac{1}{\gamma+1}\sum_{k=1}^{n}\pi_k\left(1+\gamma\lambda'\left(\boldsymbol{R}_k-\boldsymbol{R}_f\right)\right)^{\left(\frac{\gamma+1}{\gamma}\right)} I_{\Lambda_\gamma(\boldsymbol{R}_k)}(\lambda), \tag{1.59}$$

ii) if $\gamma < 0$:

$$\lambda_\gamma^* = \underset{\lambda\in \Lambda_\gamma}{arg\ max} -\frac{1}{\gamma+1}\sum_{k=1}^{n}\pi_k\left(1+\gamma\lambda'\left(\boldsymbol{R}_k-\boldsymbol{R}_f\right)\right)^{\left(\frac{\gamma+1}{\gamma}\right)}, \tag{1.60}$$

iii) if $\gamma = 0$, the maximization is unconstrained:

$$\lambda_0^* = \underset{\lambda\in R^K}{arg\ max}\ -\sum_{k=1}^{n}\pi_k\, e^{\lambda'\left(\boldsymbol{R}_k-\boldsymbol{R}_f\right)}, \tag{1.61}$$

where $\Lambda_\gamma = \{\lambda \in \mathbb{R}^K : for\ k = 1,...,n,\ (1+\gamma\lambda'(\boldsymbol{R}_k-\boldsymbol{R}_f)) > 0\}$ and $\Lambda_\gamma(\boldsymbol{R}_k) = \{\lambda \in \mathbb{R}^K : (1+\gamma\lambda'(\boldsymbol{R}_k-\boldsymbol{R}_f)) > 0\}$.

The minimum dispersion risk-neutral measure can then be recovered via the first derivative of the convex conjugate, or, equivalently, from the first-order conditions of

[51] This is an optimal nonparametric estimator for $\mathbb{P}$. For more details, see Kitamura (2006).

(1.59), (1.60) and (1.61) with respect to λ, evaluated at λ^*_γ:

$$\pi_k^{\mathbb{Q}*}(\gamma,\boldsymbol{R}) = \frac{(1+\gamma\lambda_\gamma^{*\prime}(\boldsymbol{R}_k-\boldsymbol{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\boldsymbol{R}_k)}(\lambda^*_\gamma)}{\sum_{i=1}^n (1+\gamma\lambda_\gamma^{*\prime}(\boldsymbol{R}_i-\boldsymbol{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\boldsymbol{R}_i)}(\lambda^*_\gamma)}\,,\ k=1,...,n;\ \gamma>0, \tag{1.62}$$

$$\pi_k^{\mathbb{Q}*}(\gamma,\boldsymbol{R}) = \frac{\left(1+\gamma\lambda_\gamma^{*\prime}\left(\boldsymbol{R}_k-\boldsymbol{R}_f\right)\right)^{\frac{1}{\gamma}}}{\sum_{i=1}^n \left(1+\gamma\lambda_\gamma^{*\prime}\left(\boldsymbol{R}_i-\boldsymbol{R}_f\right)\right)^{\frac{1}{\gamma}}},\ k=1,...,n;\ \gamma<0, \tag{1.63}$$

$$\pi_k^{\mathbb{Q}*}(0,\boldsymbol{R}) = \frac{e^{\lambda_0^*\left(\boldsymbol{R}_k-\boldsymbol{R}_f\right)}}{\sum_{i=1}^n e^{\lambda_0^*\left(\boldsymbol{R}_i-\boldsymbol{R}_f\right)}},\ k=1,...,n;\ \gamma=0. \tag{1.64}$$

Proof. The sample version is a particular case of the population version, so the duality result and the dual problem expressions (1.59), (1.60) and (1.61) follow directly from Theorem 1 and Corollary 2. Moreover, note that the $\delta(.)$ function has been eliminated. This is because in the sample problem, the region $\Lambda_\gamma(\mathbf{R})$ where we search for λ is simplified to depend on all observed returns at once (as in Λ_γ), instead of being a random region depending on each possible realization of returns $\mathbf{R}(\omega)$. Therefore, for $\gamma<0$, we can directly constrain the maximization by restricting λ to Λ_γ.

The minimum discrepancy risk-neutral measure can be recovered according to Corollary 3. Note that, in the sample version, we have for each realization k of the nature:

$$\frac{\pi_k^{\mathbb{Q}*}}{\pi_k} = (1+\gamma\lambda_\gamma^{*\prime}(\mathbf{R}_k-\mathbf{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\mathbf{R}_k)}(\lambda^*_\gamma)\,,\ k=1,...,n;\ \gamma>0, \tag{1.65}$$

which can be written as $\pi_k^{\mathbb{Q}*}=\pi_k(1+\gamma\lambda_\gamma^{*\prime}(\mathbf{R}_k-\mathbf{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\mathbf{R}_k)}(\lambda^*_\gamma)$. Note also that, since $\sum_{k=1}^n \pi_k(\pi_k^{\mathbb{Q}*}/\pi_k)=1 \Rightarrow \sum_{k=1}^n \pi_k^{\mathbb{Q}*}=1 \Rightarrow \sum_{k=1}^n \pi_k(1+\gamma\lambda_\gamma^{*\prime}(\mathbf{R}_k-\mathbf{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\mathbf{R}_k)}(\lambda^*_\gamma)=1$, we can write:

$$\frac{\pi_k^{\mathbb{Q}*}}{\pi_k} = \frac{(1+\gamma\lambda_\gamma^{*\prime}(\mathbf{R}_k-\mathbf{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\mathbf{R}_k)}(\lambda^*_\gamma)}{\sum_{i=1}^n \pi_i(1+\gamma\lambda_\gamma^{*\prime}(\mathbf{R}_i-\mathbf{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\mathbf{R}_i)}(\lambda^*_\gamma)}\,,\ k=1,...,n;\ \gamma>0, \tag{1.66}$$

which implies that:

$$\pi_k^{\mathbb{Q}*} = \frac{(1+\gamma\lambda_\gamma^{*\prime}(\mathbf{R}_k-\mathbf{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\mathbf{R}_k)}(\lambda^*_\gamma)}{\sum_{i=1}^n (1+\gamma\lambda_\gamma^{*\prime}(\mathbf{R}_i-\mathbf{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\mathbf{R}_i)}(\lambda^*_\gamma)}\,,\ k=1,...,n;\ \gamma>0. \tag{1.67}$$

Following the same argument, we obtain expressions (1.63) and (1.64). That is, the sample version allows us to directly identify the risk-neutral measure weights. $\square$

1.12.2 Theoretical Results and Proofs

Proof of Equivalence Between Dual Problem and Optimal Portfolio Problem

Proposition 5. *Consider the class of HARA utility functions:*

$$u^{\gamma}(W) = -\frac{1}{\gamma+1}(b - a\gamma W)^{\frac{\gamma+1}{\gamma}}, \tag{1.68}$$

with $a > 0$ and $b - a\gamma W > 0$, which guarantees that the function u^{γ} is well-defined, concave and strictly increasing. Now suppose a standard model of optimal portfolio choice, where an investor distributes her initial wealth W_0 putting $\tilde{\lambda}_j$ units of wealth on the risky asset R_j and the remaining $W_0 - \sum_{j=1}^{K} \tilde{\lambda}_j$ in a risk-free asset paying R_f. Terminal wealth is then given by $W(\tilde{\lambda}) = W_0 R_f + \sum_{j=1}^{K} \tilde{\lambda}_j (R_j - R_f)$. The investor maximizes the expected HARA utility function in (1.68) solving one of the following optimal portfolio problems:

$$\max_{\tilde{\lambda} \in R^K} \mathbb{E}\left[u^{\gamma}(W(\tilde{\lambda})) I_{\Lambda_{\gamma}(W)}(\tilde{\lambda})\right], \ \gamma > 0, \tag{1.69}$$

$$\max_{\tilde{\lambda} \in R^K} \mathbb{E}\left[u^{\gamma}(W(\tilde{\lambda})) - \delta(\tilde{\lambda} \,|\, \Lambda_{\gamma}(W))\right], \ \gamma < 0, \tag{1.70}$$

$$\max_{\tilde{\lambda} \in R^K} \mathbb{E}\left[u^{0}(W(\tilde{\lambda}))\right], \ \gamma = 0, \tag{1.71}$$

where $\Lambda_{\gamma}(W) = \{\lambda \in \mathbb{R}^K : b - a\gamma W(\tilde{\lambda}) > 0\}$. Then, solving (1.69), (1.70) and (1.71) is equivalent to solving (1.39), (1.40), and (1.41), respectively. In particular, letting λ_{γ}^{*} denote the dual problem solution, $\tilde{\lambda}_{\gamma}^{*} = -\lambda_{\gamma}^{*}(b - a\gamma W_0 R_f)/a$ if $\gamma \neq 0$ and $\tilde{\lambda}_{\gamma}^{*} = -\lambda_{\gamma}^{*}/a$ if $\gamma = 0$.

Proof. Substituting (1.68) in (1.69), we have:

$$\max_{\lambda \in R^K} \ -\frac{1}{\gamma+1}\mathbb{E}\left[(b - a\gamma W_0 R_f - a\gamma \tilde{\lambda}'(\mathbf{R} - \mathbf{R}_f))^{\frac{\gamma+1}{\gamma}} \mathbb{1}_{\{b - a\gamma W > 0\}}\right], \tag{1.72}$$

which can be written as:

$$\max_{\lambda \in R^K} \ -\frac{1}{\gamma+1}\mathbb{E}\left[((b - a\gamma W_0 R_f)(1 + \gamma \lambda'(\mathbf{R} - \mathbf{R}_f))^{\frac{\gamma+1}{\gamma}} \mathbb{1}_{\{(b - a\gamma W_0 R_f)(1 + \gamma\lambda'(\mathbf{R} - \mathbf{R}_f)) > 0\}}\right], \tag{1.73}$$

where $\lambda = -\tilde{\lambda} a/(b - a\gamma W_0 R_f)$. Note that $b - a\gamma W > 0 \Rightarrow b - a\gamma W_0 R_f > 0$, which implies that this term does not affect the indicator function. Furthermore, we can put this term in evidence outside the expectation (with exponent $\gamma/(\gamma+1)$). Since it is positive, it does not affect the maximization, so we can ignore it. This implies that (1.73) is equivalent to:

$$\max_{\lambda \in R^K} \ -\frac{1}{\gamma+1}\mathbb{E}\left[(1 + \gamma \lambda'(\mathbf{R} - \mathbf{R}_f))^{\frac{\gamma+1}{\gamma}} \mathbb{1}_{\{1 + \gamma\lambda'(\mathbf{R} - \mathbf{R}_f > 0\}}\right], \tag{1.74}$$

which is precisely the dual problem (1.39).

The argument is analogous for the equivalence between (1.69) and the dual problem (1.40).

In the case that $\gamma = 0$, we take the limit $\gamma \to 0$ of (1.71) to get:

$$\underset{\lambda \in R^K}{max} \; -\mathbb{E}\left[e^{-aW}\right] \equiv \underset{\lambda \in R^K}{max} \; -\mathbb{E}\left[e^{-aW_0R_f - a\tilde{\lambda}'(\mathbf{R}-\mathbf{R}_f)}\right]. \tag{1.75}$$

By taking $e^{-aW_0R_f}$ outside of the expectation, note that it does not affect the maximization, so we can ignore it. By letting $\lambda = -a\tilde{\lambda}$, we have the desired result. □

Proof of Proposition 1 in Main Paper

Proposition 6. *Consider the minimum dispersion primal problem (1.58) and the corresponding dual problems in the usual case where there is risk premia in the economy. Then, neither the primal nor the dual problems have a solution when $\gamma \to -\infty$ or $\gamma \to \infty$.*

Proof. Consider first the primal problem (1.58). One of its constraints is that $\sum_{k=1}^{n} \pi_k^{\mathbb{Q}} = 1$, which implies that $\frac{1}{n}\sum_{k=1}^{n} \pi_k^{\mathbb{Q}} = \frac{1}{n} = \pi_k$, that is, on average the risk-neutral measure equals π_k. Therefore, since $\pi_k^{\mathbb{Q}} \neq \pi_k$ for all k, $\pi_k^{\mathbb{Q}}/\pi_k$ will always have values smaller and larger than 1. This implies that when $\gamma \to -\infty$ ($\gamma \to \infty$), the values smaller than 1 (larger than 1) will diverge to infinity, faster than the denominator, in the objective function. Therefore, in the limits the primal problem is not well-defined and does not have a solution.

Consider now the dual problem for $\gamma < 0$. If $\gamma \to -\infty$, we have $(1 + \gamma\lambda_\gamma^{*\prime}(\mathbf{R}_k - \mathbf{R}_f))^{\frac{1}{\gamma}} \to 1$, so the risk-neutral measure converges to the physical measure: $\pi_k^{\mathbb{Q}*} \to \pi_k = 1/n$. Since there is risk premia in the economy, this measure does not price the basis assets, hence not satisfying the first-order conditions of the dual problem, implying that the dual problem does not have a solution.

Now, consider the dual problem for $\gamma > 0$. If $\gamma \to \infty$, we have $(1 + \gamma\lambda_\gamma^{*\prime}(\mathbf{R}_k - \mathbf{R}_f))^{\frac{1}{\gamma}} \to 1$. Let's analyze the behavior of the indicator function in the limit. For states of nature k such that $\lambda_\gamma^{*\prime}(\mathbf{R}_k - \mathbf{R}_f) \geq 0$ in the limit, the indicator function will be one, otherwise it will be zero. This can be seen by rewriting, for a finite γ, the condition for the indicator function to be one as $\lambda_\gamma^{*\prime}(\mathbf{R}_k - \mathbf{R}_f) > -1/\gamma$, and noticing that in the limit it becomes $\lambda_\gamma^{*\prime}(\mathbf{R}_k - \mathbf{R}_f) \geq 0$. Therefore, the limit of the risk-neutral measure when $\gamma \to \infty$ is $1/\beta$ for states of nature such that $\lambda_\gamma^{*\prime}(\mathbf{R}_k - \mathbf{R}_f) \geq 0$, where β is the number of such states, and 0 otherwise. This measure does not price the basis assets, implying that the dual problem does not have a solution.

Proof of Proposition 2 in Main Paper

Proposition 7. *Let λ_γ^* be the solution for a given γ to the dual problem in Corollary 4 when there is just one risky basis asset.*

(i) Let $\gamma < 0$ in a neighborhood of zero. If $\mathbb{E}(R - R_f) > 0$ ($\mathbb{E}(R - R_f) < 0$), we have $\lambda_\gamma^* < 0$ ($\lambda_\gamma^* > 0$) and $\gamma\lambda_\gamma^* > 0$ ($\gamma\lambda_\gamma^* < 0$). As γ becomes more negative, $\gamma\lambda_\gamma^*$ increases (decreases). The dual problem has a solution while $\gamma\lambda_\gamma^* < -1/min\{R_k - R_f\}$ ($\gamma\lambda_\gamma^* > -1/max\{R_k - R_f\}$). When this condition can no longer be satisfied, there is no longer a solution to the dual problem.

(ii) Let $\gamma > 0$ in a neighborhood of zero. If $\mathbb{E}(R - R_f) > 0$ ($\mathbb{E}(R - R_f) < 0$), we have $\lambda_\gamma^* < 0$ ($\lambda_\gamma^* > 0$) and $\gamma\lambda_\gamma^* < 0$ ($\gamma\lambda_\gamma^* > 0$). As γ increases, $\gamma\lambda_\gamma^*$ decreases (increases). While $\gamma\lambda_\gamma^* > -1/max\{R_k - R_f\}$ ($\gamma\lambda_\gamma^* < -1/min\{R_k - R_f\}$), the implied risk-neutral measure is strictly positive. When $\gamma\lambda_\gamma^* \leq -1/max\{R_k - R_f\}$ ($\gamma\lambda_\gamma^* \geq -1/min\{R_k - R_f\}$), the measure is not strictly positive anymore. Let also max_i (min_i) denote the i^{th} highest (smallest) value and $min_{>0}$ ($max_{<0}$) the minimum positive (maximum negative) value. The i^{th} zero in the measure will appear when $\gamma\lambda_\gamma^* \leq -1/max_i\{R_k - R_f\}$ ($\gamma\lambda_\gamma^* \geq -1/min_i\{R_k - R_f\}$). When the last positive (negative) return is set to zero, i.e., $\gamma\lambda_\gamma^* \leq -1/min_{>0}\{R_k - R_f\}$ ($\gamma\lambda_\gamma^* \geq -1/max_{<0}\{R_k - R_f\}$), the dual problem does not have a solution anymore.

Proof. It is a well known result that, for the portfolio optimization problem in Proposition 5 with one risky asset and any utility function with $u'(.) > 0$ and $u''(.) < 0$, we have, denoting the solution $\tilde{\lambda}^*$:

$$\tilde{\lambda}^* > 0 \iff \mathbb{E}(R - R_f) > 0,$$

$$\tilde{\lambda}^* = 0 \iff \mathbb{E}(R - R_f) = 0,$$

$$\tilde{\lambda}^* < 0 \iff \mathbb{E}(R - R_f) < 0.$$

In general, $\mathbb{E}(R - R_f) > 0$ and this implies that the investor will always hold a positive amount of the asset. However, it is possible that $\mathbb{E}(R - R_f) < 0$, in the case of assets indexed on catastrophe events, for instance. Therefore, we prove our results considering both possibilities. Proposition 1 shows that, denoting λ^* to be the dual problem solution, we have that $\tilde{\lambda}^* > 0$ implies $\lambda^* < 0$ and that $\tilde{\lambda}^* < 0$ implies $\lambda^* > 0$.

In order to study the solutions of the dual problem, note that it is strictly concave, implying that there exists a solution if and only if the first-order condition (FOC) with respect to λ is satisfied:

$$\sum_{k=1}^{n} \left(1 + \gamma\lambda_\gamma^*(R_k - R_f)\right)^{\frac{1}{\gamma}} (R_k - R_f) = 0, \text{ if } \gamma < 0, \tag{1.76}$$

$$\sum_{k=1}^{n} \left(1+\gamma\lambda_\gamma^*(R_k-R_f)\right)^{\frac{1}{\gamma}} I_{\Lambda_\gamma(R_k)}(\lambda_\gamma^*)(R_k-R_f)=0, \text{ if } \gamma>0. \tag{1.77}$$

In addition, a solution for the dual problem when $\gamma < 0$ must also satisfy Λ_γ. Note also that the FOC imposes that the solution of the dual problem satisfies the risk neutrality constraint of the primal problem.

(i) Let $\gamma < 0$. We are going to analyze how the expression given by the FOC (1.76) changes with respect to γ and $\gamma\lambda_\gamma^*$. By the implicit function theorem, given a solution (γ^*, λ^*) it is always possible to obtain a λ satisfying the FOC in an open interval containing γ^*. Whenever γ varies, $\gamma\lambda_\gamma^*$ will change in response in order to continue to set the FOC equal to zero. Therefore, we can look at the variations of the FOC with respect to $\gamma\lambda_\gamma^*$ and γ separately. Being so, let $u \equiv \gamma\lambda_\gamma^*$. Recall that when $\gamma < 0$, if $\lambda_\gamma^* < 0$ ($\lambda_\gamma^* > 0$), we have $u > 0$ ($u < 0$). Rewriting the FOC:

$$\sum_{k=1}^{n} \left(1+u(R_k-R_f)\right)^{\frac{1}{\gamma}} (R_k-R_f)=0.$$

Start from a negative γ close to zero for which the dual problem has a solution by the implicit function theorem, since for $\gamma \to 0$ the problem is unconstrained and there is a solution. Then, $\sum_{k=1}^{n} \left(1+u(R_k-R_f)\right)^{\frac{1}{\gamma}} (R_k-R_f)=0$ and u satisfies Λ_γ. Taking the derivative of the expression in the FOC with respect to γ, if $\lambda_\gamma^* < 0$ ($\lambda_\gamma^* > 0$), we have:

$$-\sum_{k=1}^{n} \frac{1}{\gamma^2} \left(1+u(R_k-R_f)\right)^{\frac{1}{\gamma}} \ln(1+u(R_k-R_f))(R_k-R_f) \ < 0 \quad (>0),$$

because $(1+u(R_k-R_f)) > 0$, since it is a solution and u satisfies Λ_γ, and $\ln(1+u(R_k-R_f))$ has the same sign as (R_k-R_f) if $u > 0$ (the opposite sign if $u < 0$), so that the product inside the sum for each k is positive (negative). This implies that when we vary γ to be more negative, the FOC increases if $\lambda_\gamma^* < 0$, and decreases if $\lambda_\gamma^* > 0$. Now we take the derivative of the FOC with respect to u:

$$\sum_{k=1}^{n} \frac{1}{\gamma} \left(1+u(R_k-R_f)\right)^{\frac{1}{\gamma}-1} (R_k-R_f)^2 < 0.$$

That is, when u increases (decreases), the FOC decreases (increases). Therefore, when we are at a solution of the dual problem and we make γ slightly more negative, if $\lambda_\gamma^* < 0$ ($\lambda_\gamma^* > 0$), the FOC increases (decreases) and in order to continue having a solution u must increase (decrease) to compensate. Remember that a solution of the dual problem must also satisfy:

$$1 + u(R_k - R_f) > 0 \;\; \forall \; k \Leftrightarrow 1 + u \; \min\{R_k - R_f\} > 0 \Leftrightarrow u < \frac{-1}{\min\{R_k - R_f\}}, \;\; \text{if} \;\; u > 0,$$

$$1 + u(R_k - R_f) > 0 \;\; \forall \; k \Leftrightarrow 1 + u \; \max\{R_k - R_f\} > 0 \Leftrightarrow u > \frac{-1}{\max\{R_k - R_f\}}, \;\; \text{if} \;\; u < 0.$$

Thus, if $u > 0$ ($u < 0$), there is an upper (lower) bound for u to be a solution. Therefore, the problem can not compensate a smaller γ with a larger u indefinitely, in the case that $u > 0$, or with a smaller u indefinitely if $u < 0$. The dual problem breaks down when we can not make u larger or smaller in order to continue to have a solution. More specifically, if $\lambda^*_\gamma < 0$ ($\lambda^*_\gamma > 0$), when $u \geq \frac{-1}{\min\{R_k - R_f\}}$ ($u \leq \frac{-1}{\max\{R_k - R_f\}}$) the dual problem does not have a solution anymore.

(ii) When $\gamma > 0$, if $\lambda^*_\gamma < 0$ ($\lambda^*_\gamma > 0$), we have $u < 0$ ($u > 0$). Rewriting the FOC (1.77):

$$\sum_{k=1}^{n} (1 + u(R_k - R_f))^{\frac{1}{\gamma}} \, I_{\Lambda_\gamma(R_k)}(\lambda^*_\gamma)(R_k - R_f) = 0.$$

Start from a positive γ close to zero for which the dual problem has a solution by the implicit function theorem, since for $\gamma \to 0$ the problem is unconstrained and there is a solution. Then, $\sum_{k=1}^{n} (1 + u(R_k - R_f))^{\frac{1}{\gamma}} \, I_{\Lambda_\gamma(R_k)}(\lambda^*_\gamma)(R_k - R_f) = 0$. Taking the derivative of the expression in the FOC with respect to γ, if $\lambda^*_\gamma < 0$ ($\lambda^*_\gamma > 0$), we have:

$$-\sum_{k=1}^{n} \frac{1}{\gamma^2} (1 + u(R_k - R_f))^{\frac{1}{\gamma}} \, I_{\Lambda_\gamma(R_k)}(\lambda^*_\gamma) \ln(1 + u(R_k - R_f))(R_k - R_f) \; > 0 \quad (< 0),$$

because $(1 + u(R_k - R_f)) > 0$, otherwise the indicator function would be set to zero, and $\ln(1 + u(R_k - R_f))$ has the same sign as $(R_k - R_f)$ if $u > 0$ (the opposite sign if $u < 0$), so that the product inside the sum for each k is positive (negative) or zero if there is an active indicator function. This implies that when we increase γ, the FOC increases if $\lambda^*_\gamma < 0$, and decreases if $\lambda^*_\gamma > 0$. Now we take the derivative of the FOC with respect to u:

$$-\sum_{k=1}^{n} \frac{1}{\gamma} (1 + u(R_k - R_f))^{\frac{1}{\gamma}-1} \, I_{\Lambda_\gamma(R_k)}(\lambda^*_\gamma)(R_k - R_f)^2 < 0.$$

That is, when u increases (decreases), the FOC decreases (increases). Therefore, when we are at a solution of the dual problem and we make γ slightly larger, if $\lambda^*_\gamma < 0$ ($\lambda^*_\gamma > 0$), the FOC increases (decreases) and in order to continue having a solution u must decrease (increase) to compensate. As we increase γ, the indicator functions will all be one until the following is satisfied:

$$1 + u(R_k - R_f) > 0 \;\; \forall\, k \Leftrightarrow 1 + u \, \max\{R_k - R_f\} > 0 \Leftrightarrow u > \frac{-1}{\max\{R_k - R_f\}}, \;\; \text{if } \; u < 0,$$

$$1 + u(R_k - R_f) > 0 \;\; \forall\, k \Leftrightarrow 1 + u \, \min\{R_k - R_f\} > 0 \Leftrightarrow u < \frac{-1}{\min\{R_k - R_f\}}, \;\; \text{if } \; u > 0.$$

Thus, if $u < 0$ ($u > 0$), there is a lower (upper) bound for u so that all indicator functions are active. If $\lambda_\gamma^* < 0$ ($\lambda_\gamma^* > 0$), when $u \leq \frac{-1}{\max\{R_k - R_f\}}$ ($u \geq \frac{-1}{\min\{R_k - R_f\}}$), the risk-neutral measure sets the maximum (minimum) return to zero. From that point on, the implied measure is not strictly positive anymore. However, the dual problem will continue to have a solution, because it is unconstrained and the measures will be able to continue to zero out the positive (negative) returns in order to satisfy the FOC. Let $\max_i$ and $\min_{>0}$ ($\min_i$ and $\max_{<0}$) denote the i^{th} highest (smallest) value and the minimum positive (maximum negative) value, respectively. The second zero in the measure will appear when $u \leq \frac{-1}{\max_2\{R_k - R_f\}}$ ($u \geq \frac{-1}{\min_2\{R_k - R_f\}}$), the third zero when $u \leq \frac{-1}{\max_3\{R_k - R_f\}}$ ($u \geq \frac{-1}{\min_3\{R_k - R_f\}}$), and so on. This will continue in this fashion until the last positive (negative) return is set to zero, i.e., when $u \leq \frac{-1}{\min_{>0}\{R_k - R_f\}}$ ($u \geq \frac{-1}{\max_{<0}\{R_k - R_f\}}$). From this point on, the dual problem does not have a solution anymore, because the measure puts weights only on negative (positive) returns. $\square$

Constructing Minimum Dispersion Price Bounds from K Basis Assets

In the main paper, we focus on obtaining option price bounds from the returns of a single risky asset (the underlying). In this subsection, we show how to construct minimum dispersion price bounds for any non-redundant asset from returns on K basis assets. The next proposition provides guidance on how to identify the set of admissible minimum dispersion measures in this case.

Proposition 8. *Consider the dual problem as in Corollary 4.*

(i) Let $\gamma < 0$. Restriction Λ_γ is equivalent to $0 <$ $max_k\{\lambda_\gamma^{*\prime}(\boldsymbol{R}_k - \boldsymbol{R}_f)\} < -1/\gamma$. The dual problem has a solution while this restriction is satisfied. As $\gamma \to -\infty$, the dual problem does not satisfy the constraint and does not have a solution anymore.

(ii) Let $\gamma > 0$. The indicator function $I_{\Lambda_\gamma(\boldsymbol{R}_k)}$ is equal to one for all k if $0 >$ $min_k\{\lambda_\gamma^{*\prime}(\boldsymbol{R}_k - \boldsymbol{R}_f)\} > -1/\gamma$. A strictly positive solution is guaranteed while this restriction is satisfied. As $\gamma \to \infty$, the implied risk-neutral measure will not be strictly positive anymore.

Proof. Consider the first-order conditions for the dual problem for $\gamma < 0$ (it will be analogous for $\gamma > 0$):

$$\sum_{k=1}^{n}(1+\gamma\lambda_{\gamma}^{*\prime}(\mathbf{R}_k-\mathbf{R}_f))^{\frac{1}{\gamma}}(R_k^j-R_f)=0,\;\; j=1,...,K. \tag{1.78}$$

This implies that the risk-neutral measure must price all the basis assets and any portfolio based on the basis assets. In particular, it must price the optimal portfolio:

$$\sum_{k=1}^{n}(1+\gamma\lambda_{\gamma}^{*\prime}(\mathbf{R}_k-\mathbf{R}_f))^{\frac{1}{\gamma}}(\lambda_{\gamma}^{*\prime}(\mathbf{R}_k-\mathbf{R}_f))=0. \tag{1.79}$$

The condition above has no solution if $\lambda_{\gamma}^{*\prime}(\mathbf{R}_k-\mathbf{R}_f)$ does not alternate in sign, which would imply that the portfolio has a return larger than the risk-free asset in all states of nature, providing an arbitrage. Therefore, it has to alternate signs. This implies that $\max_k\{\lambda_{\gamma}^{*\prime}(\mathbf{R}_k-\mathbf{R}_f)\}>0$ and $\min_k\{\lambda_{\gamma}^{*\prime}(\mathbf{R}_k-\mathbf{R}_f)\}<0$. Note also that $\lambda_{\gamma}^{*}\neq\mathbf{0}$, because otherwise the risk-neutral measure would be equal to 1 in all states of nature, not pricing the basis assets in the presence of risk premia in the economy.

(i) Consider the case with $\gamma<0$. A solution to the dual problem satisfies:

$$1+\gamma\lambda_{\gamma}^{*\prime}(\mathbf{R}_k-\mathbf{R}_f)>0\;\;\forall\, k, \tag{1.80}$$

which is equivalent to:

$$1+\gamma\max_k\{\lambda_{\gamma}^{*\prime}(\mathbf{R}_k-\mathbf{R}_f)\}>0, \tag{1.81}$$

because $\gamma<0$ and $\max_k\{\lambda_{\gamma}^{*\prime}(\mathbf{R}_k-\mathbf{R}_f)\}>0$, making $\gamma\max_k\{\lambda_{\gamma}^{*\prime}(\mathbf{R}_k-\mathbf{R}_f)\}<0$ the minimum value across the n states of nature. Combining (1.81) with the fact that $\max_k\{\lambda_{\gamma}^{*\prime}(\mathbf{R}_k-\mathbf{R}_f)\}>0$, we have:

$$0<\max_k\{\lambda_{\gamma}^{*\prime}(\mathbf{R}_k-\mathbf{R}_f)\}<-\frac{1}{\gamma}. \tag{1.82}$$

For $\gamma\to 0$, the problem is unconstrained and there will be a solution. When $\gamma\to-\infty$, condition (1.82) breaks down, and the dual problem does not have a solution. It remains to show that, as γ decreases, there will be a solution while (1.82) is satisfied. By the maximum theorem, the function $h(\gamma)$ giving the *argmax* of the dual problem is continuous. Moreover, by the implicit function theorem, given a solution (γ^*,λ^*) it is always possible to obtain a λ satisfying the first-order condition in an open interval containing γ^*. Therefore, as γ decreases, λ can change to continue satisfying the first-order condition, where $h(\gamma)$ changes continuously until it is not possible to satisfy condition (1.82).

(ii) Consider now that $\gamma > 0$. The dual problem is unconstrained, but contains an indicator function that sets to zero weights to states of nature where $1+\gamma\lambda_\gamma^{*\prime}(\mathbf{R}_k-\mathbf{R}_f) \leq 0$. Notice that $1+\gamma \min_k\{\lambda_\gamma^{*\prime}(\mathbf{R}_k-\mathbf{R}_f)\} > 0$ implies that $1+\gamma\lambda_\gamma^{*\prime}(\mathbf{R}_k-\mathbf{R}_f) > 0$ for all k, because since $\gamma > 0$, $\gamma \min_k\{\lambda_\gamma^{*\prime}(\mathbf{R}_k-\mathbf{R}_f)\} < 0$ is the minimum value across states of nature. Therefore, there will be no zero weights as long as the following holds:

$$0 > \min_k\{\lambda_\gamma^{*\prime}(\mathbf{R}_k-\mathbf{R}_f)\} > -\frac{1}{\gamma}. \tag{1.83}$$

For $\gamma \to 0$, (1.83) will hold. For $\gamma \to \infty$, condition (1.83) breaks down, and the risk-neutral measure is not strictly positive. Starting from γ close to zero, as γ increases, the implicit function theorem guarantees strictly positive solutions of the dual problem while (1.83) is satisfied, where the first-order condition is differentiable. When $\min_k\{\lambda_\gamma^{*\prime}(\mathbf{R}_k-\mathbf{R}_f)\} \leq -1/\gamma$, the guaranteed solutions will not be strictly positive anymore. $\square$

The proposition above suggests an algorithm to find the set of admissible minimum dispersion risk-neutral measures. In practice, the results will kick in way before γ approaches minus or plus infinity, where, as γ becomes more negative, the constraint will be satisfied until some negative γ, while as γ becomes larger, the measures will be strictly positive until a given γ is reached. Therefore, one can solve the dual problem for a grid of γ's becoming more negative and more positive, while calculating $\max_k\{\lambda_\gamma^{*}(\mathbf{R}_k-\mathbf{R}_f)\}$ and $\min_k\{\lambda_\gamma^{*}(\mathbf{R}_k-\mathbf{R}_f)\}$. The negative $\underline{\gamma}$ where $\max_k\{\lambda_\gamma^{*}(\mathbf{R}_k-\mathbf{R}_f)\}$ approaches $-1/\gamma$ and positive $\overline{\gamma}$ where $\min_k\{\lambda_\gamma^{*}(\mathbf{R}_k-\mathbf{R}_f)\} \leq -1/\gamma$ define the interval $[\underline{\gamma},\overline{\gamma}]$ of strictly positive minimum dispersion risk-neutral measures.[52] This set is endogenously determined by the basis assets returns. It is also possible to include measures with zeros in some states of nature by allowing for γ's greater than $\overline{\gamma}$.

Since $[\underline{\gamma},\overline{\gamma}]$ explicitly identifies all the measures within the set, it is possible to calculate the implied price of a non-redundant asset by each measure, given a grid. Moreover, because the risk-neutral measures are continuous functions of γ, the implied prices will change continuously with γ and reach a maximum and minimum in the interval. One can then calculate upper and lower price bounds for the non-redundant asset by taking the maximum and minimum prices implied by the measures in $[\underline{\gamma},\overline{\gamma}]$. More specifically, let x be the non-redundant payoff we want to price. The lower minimum dispersion price

[52] For $\gamma < 0$, the problem is constrained by Λ_γ. Computationally, this constraint is imposed and will not be violated. Instead, the dual problem stops having a solution in practice when $\max_k\{\lambda_\gamma^{*}(\mathbf{R}_k-\mathbf{R}_f)\}$ gets arbitrarily close to $-1/\gamma$, up to the fifth decimal place. Solutions beyond this point will be associated to non-negligible pricing errors for the basis assets, which is the computational counterpart of the dual problem not having a solution anymore. For $\gamma > 0$, $\min_k\{\lambda_\gamma^{*}(\mathbf{R}_k-\mathbf{R}_f)\} > -1/\gamma$ is not a constraint to the problem, so it will become $\min_k\{\lambda_\gamma^{*}(\mathbf{R}_k-\mathbf{R}_f)\} \leq -1/\gamma$ at some point. An appropriate grid of γ's is with 0.1 spacing. In general, the γ's defining the pricing interval will be between -20 and 20.

bound solves:

$$\underline{C} = \min_{\{\pi_\gamma^Q\}} \frac{1}{R_f} \sum_{k=1}^{n} \pi_{\gamma k}^{\mathbb{Q}}\, x_k, \;\; s.t. \;\; \gamma \in [\underline{\gamma}, \overline{\gamma}], \tag{1.84}$$

where the upper bound $\overline{C}$ solves the corresponding maximum. If, *a posteriori*, the asset price C is observed and $C \in [\underline{C}, \overline{C}]$, one can recover the risk-neutral measure and the implied γ that correctly price the asset.

Proof of Proposition 3 in Main Paper

Proposition 9. *Suppose a Black-Scholes economy with drift* μ*, risk-free rate* r *and volatility* σ*. Then, there is an optimal Cressie-Read discrepancy indexed by* γ^* *for which the implied minimum dispersion option price equals the Black-Scholes price, given by:*

$$\gamma^* = -\frac{\sigma^2}{\mu - r}\ . \tag{1.85}$$

Proof. Rubinstein (1976) proves that if the underlying stock price and the SDF are jointly lognormal, the Black-Scholes option pricing formula obtains without dynamic trading. A sufficient condition for the lognormality condition is that the SDF comes from a CRRA utility function, i.e., $m = R^{-\eta}$ with $\eta > 0$, where R is the compound gross rate of return of the underlying asset through the expiration t of the option. The parameter η can be obtained from the stochastic process of R. From the pricing equation $1 = \mathbb{E}[mR]$, where m is an admissible SDF, we have $1 = cov(R, m) + \mathbb{E}[R]\mathbb{E}[m]$. Since $\mathbb{E}[m] = 1/R_f$, we can write this equation as:

$$R_f = \frac{Cov(R, m)}{\mathbb{E}[m]} + \mathbb{E}[R], \tag{1.86}$$

where $R_f = (1+r)^t$, $r = e^{r_c} - 1$ and r_c is the continuous time risk-free rate of the model. Substituting $m = R^{-\eta}$, we have:

$$(1+r)^t = \frac{Cov(R, R^{-\eta})}{\mathbb{E}[R^{-\eta}]} + \mathbb{E}[R], \tag{1.87}$$

$$\Rightarrow (1+r)^t = \frac{\mathbb{E}[R^{1-\eta}]}{\mathbb{E}[R^{-\eta}]}. \tag{1.88}$$

Under lognormality of R, $R^{1-\eta}$ and $R^{-\eta}$ are lognormal, implying that:

$$(1+r)^t = \frac{exp[(1-\eta)\mathbb{E}(\ln R) + \frac{1}{2}(1-\eta)^2 Var(\ln R)]}{exp[(-\eta)\mathbb{E}(\ln R) + \frac{1}{2}\eta^2 Var(\ln R)]}, \tag{1.89}$$

which can be simplified to:

$$\eta = \frac{\mathbb{E}(\ln R) - t\ln(1+r)}{Var(\ln R)} + \frac{1}{2}. \tag{1.90}$$

Under the Black-Scholes economy, $\ln R \sim N((\mu - \sigma^2/2)t, \sigma^2 t)$. Using this in (1.90), we have:

$$\eta = \frac{\mu - r_c}{\sigma^2}. \tag{1.91}$$

From Proposition 1, when $b = 0$ and $a = -1/\gamma$, the minimum dispersion risk-neutral measure is equivalent to the one implied by CRRA utility. Comparing with $m = R^{-\eta}$, we have $\gamma = -1/\eta$. Therefore, the Cressie-Read discrepancy that generates the risk-neutral measure for which the Black-Scholes option pricing equation is valid is given by:

$$\gamma^* = -\frac{\sigma^2}{\mu - r_c}. \tag{1.92}$$

$\square$

Proof of Proposition 4 in Main Paper

We follow Cerny (2003) in measuring the attractiveness of a self-financing investment by the certainty equivalent of the resulting wealth W relative to the wealth of a riskless investment. The value of the best deal in the market with excess returns $\mathbf{R}$, denoted $\alpha(\mathbf{R})$, is defined implicitly as:

$$u^\gamma(W_0 R_f + \alpha(\mathbf{R})) \equiv \max_\lambda \mathbb{E}[u^\gamma(W_0 R_f + \lambda'(\mathbf{R} - \mathbf{R}_f))]. \tag{1.93}$$

The next proposition derives the SDF moment restrictions implied by the class of HARA utility functions.

Proposition 10. *Consider the class of HARA utility functions as in (1.12) and the environment described in Proposition 1. Then, the following SDF moment restrictions hold for $\gamma \in (-\infty, \infty)$:*[53]

$$\left(1 - \gamma A^\gamma(w_0)\alpha_{\gamma basis}\right)^{-(\gamma+1)} \leq \mathbb{E}[m^{\gamma+1}] \leq \left(1 - \gamma A^\gamma(w_0)\overline{\alpha}_\gamma\right)^{-(\gamma+1)}, \tag{1.94}$$

$$s.t.\ \mathbb{E}[m(\boldsymbol{R} - \boldsymbol{R}_f)] = \boldsymbol{0},\ \ m \geq 0,$$

where $w_0 = W_0 R_f$, $A^\gamma(w_0) = a/(b - \gamma a w_0)$, $\overline{\alpha}_\gamma \geq \alpha_{\gamma basis}$ and $\alpha_{\gamma basis}$, the certainty equivalent of the best deal attainable in the market containing only the basis assets, is given by:

$$\alpha_{\gamma basis} = \min_m\ \alpha_\gamma(m),\ \ s.t.\ \mathbb{E}[m(\boldsymbol{R} - \boldsymbol{R}_f)] = \boldsymbol{0},\ \ m \geq 0. \tag{1.95}$$

The restrictions above can be interpreted as reward-for-risk measures as follows:

$$\left(1 + h^2_{\gamma basis}\right)^{\frac{\gamma(\gamma+1)}{2}} \leq \mathbb{E}[m^{\gamma+1}] \leq \left(1 + \overline{h}^2_\gamma\right)^{\frac{\gamma(\gamma+1)}{2}}, \tag{1.96}$$

[53]The limiting cases of $\gamma \to 0$ and $\gamma \to -1$ are proved in Cerny (2003).

$$s.t. \ \mathbb{E}[m(\boldsymbol{R} - \boldsymbol{R}_f)] = \boldsymbol{0}, \ \ m \geq 0,$$

where $h_\gamma^2 = 2A^\gamma(w_0)\alpha_\gamma$ *is the generalized Sharpe ratio.*

Proof. Suppose that the market containing basis assets is complete and the state prices are given by a unique state-price density m. We aim to find the maximum certainty equivalent gain $\alpha(m)$ in this market. To that end, we search for the optimal distribution of wealth, subject to the budget constraint imposed by m:

$$\max_{\lambda} \mathbb{E}[u^\gamma(w_0 + \lambda'(\mathbf{R} - \mathbf{R}_f))] = \max_{W} \mathbb{E}[u^\gamma(W)] \ \ s.t. \ \ \mathbb{E}[mW] = w_0, \tag{1.97}$$

where $w_0 = W_0 R_f$, which implies:

$$u^\gamma(w_0 + \alpha_\gamma(m)) = \max_{W} \mathbb{E}[u^\gamma(W)] \ \ s.t. \ \ \mathbb{E}[mW] = w_0. \tag{1.98}$$

Following Cerny (2003), problem (1.98) can be solved using unconstrained maximization separately in each state with a Lagrange multiplier δ:

$$\max_{W(\omega), \omega \in \Omega} u^\gamma(W(\omega)) - \delta m(\omega) W(m). \tag{1.99}$$

The first-order conditions give $a(b - a\gamma W)^{\frac{1}{\gamma}} = \delta m$, which can be written as:

$$W = \left(-\frac{1}{a\gamma}\right)\left(\frac{\delta m}{a}\right)^\gamma + \frac{b}{a\gamma}. \tag{1.100}$$

We can substitute (1.100) in the restriction $\mathbb{E}[mW] = w_0$ to recover δ:

$$\mathbb{E}\left[m^{\gamma+1}\right]\left(\frac{\delta}{a}\right)^\gamma \left(\frac{-1}{a\gamma}\right) + \mathbb{E}[m]\frac{b}{a\gamma} = w_0,$$

$$\Rightarrow \delta^\gamma = \frac{a^\gamma(b - a\gamma w_0)}{\mathbb{E}[m^{\gamma+1}]}. \tag{1.101}$$

Substituting back in (1.100), we have:

$$W = \left(w_0 - \frac{b}{a\gamma}\right)\frac{m^\gamma}{\mathbb{E}[m^{\gamma+1}]} + \frac{b}{a\gamma}, \tag{1.102}$$

$$\Rightarrow u^\gamma(W) = -\frac{1}{\gamma+1}\left[b - a\gamma\left(w_0 - \frac{b}{a\gamma}\right)\frac{m^\gamma}{\mathbb{E}[m^{\gamma+1}]} - b\right]^{\frac{\gamma+1}{\gamma}},$$

$$\Rightarrow \mathbb{E}[u^\gamma(W)] = -\frac{1}{\gamma+1}(b - a\gamma w_0)^{\frac{\gamma+1}{\gamma}}\mathbb{E}[m^{\gamma+1}]^{-\frac{1}{\gamma}}. \tag{1.103}$$

To recover the certainty equivalent of the optimal risky investment, we use (1.103) in

(1.98):

$$-\frac{1}{\gamma+1}(b-a\gamma\alpha_\gamma-a\gamma w_0)^{\frac{\gamma+1}{\gamma}}=-\frac{1}{\gamma+1}(b-a\gamma w_0)^{\frac{\gamma+1}{\gamma}}\mathbb{E}[m^{\gamma+1}]^{-\frac{1}{\gamma}},$$

$$\alpha_\gamma(m)=\frac{(b-a\gamma w_0)}{-a\gamma}\mathbb{E}[m^{\gamma+1}]^{-\frac{1}{(\gamma+1)}}+\frac{(b-a\gamma w_0)}{a\gamma} \tag{1.104}$$

Noting that $A^\gamma(w_0)=a/(b-a\gamma w_0)$ is the absolute risk aversion in initial wealth, we have:

$$\alpha_\gamma(m)=-\frac{1}{\gamma A^\gamma(w_0)}\mathbb{E}[m^{\gamma+1}]^{-\frac{1}{(\gamma+1)}}+\frac{1}{\gamma A^\gamma(w_0)}. \tag{1.105}$$

The link between the complete and incomplete market is provided by the extension theorem, which asserts that any incomplete market without good deals can be embedded in a complete market that has no good deals. Denote by $\alpha_{\gamma basis}$ the certainty equivalent of the best deal attainable in the market containing only basis assets. Cerny (2003) shows that the extension theorem implies that:

$$\alpha_{\gamma basis}=\underset{m}{min}\ \alpha_\gamma(m), \tag{1.106}$$

where m must price correctly all basis assets. Suppose that we want to find all prices of a non-redundant asset that do not provide good deals of size $\overline{\alpha}_\gamma$ in the enlarged market. From the extension theorem, all such prices must be supported by pricing kernels for which $\alpha_\gamma(m)\leq\overline{\alpha}_\gamma$. This is the dual no-good-deal discount factor restriction. Therefore, by re-writing (1.105), the no-good-deal SDF restrictions are:

$$(1-\gamma A^\gamma(w_0)\alpha_{\gamma basis})^{-(\gamma+1)}\leq\mathbb{E}[m^{\gamma+1}]\leq(1-\gamma A^\gamma(w_0)\overline{\alpha}_\gamma)^{-(\gamma+1)}. \tag{1.107}$$

Now we interpret the restrictions (1.107) as reward-for-risk measures that are close in spirit to the Sharpe ratio. Cerny (2003) shows that there is an asymptotic relationship between the certainty equivalent α and the Sharpe ratio h: $A^\gamma(w_0)\alpha_\gamma=h_\gamma^2/2$. However, this transformation is not unique. For all κ we have:

$$(1-\gamma A^\gamma(w_0)\alpha_\gamma)^{-(\gamma+1)}=\left(1-\gamma\frac{h_\gamma^2}{2}\right)^{-(\gamma+1)}=\left(1-\kappa\gamma\frac{h_\gamma^2}{2}\right)^{-\frac{1}{\kappa}(\gamma+1)}. \tag{1.108}$$

The ambiguity is resolved by maintaining the consistency with the arbitrage-adjusted Sharpe ratio, for which the duality, as shown by Cerny (2003), is $\mathbb{E}[m^2]=1+h_A^2$. The HARA utility function obtains the truncated quadratic utility when $\gamma=1$. Therefore, by choosing $\kappa=-2/\gamma$, we obtain:

$$\left(1+h_\gamma^2\right)^{\frac{\gamma(\gamma+1)}{2}}\overset{\gamma=1}{=}1+h_1^2=1+h_A^2. \tag{1.109}$$

The restrictions in terms of generalized Sharpe ratios then become:

$$\left(1 + h^2_{\gamma basis}\right)^{\frac{\gamma(\gamma+1)}{2}} \leq \mathbb{E}[m^{\gamma+1}] \leq \left(1 + h^2_{\gamma}\right)^{\frac{\gamma(\gamma+1)}{2}}. \tag{1.110}$$

□

1.12.3 Additional Empirical Results

Summary Statistics - Berkeley Options Database

Summary statistics for the S&P 500 index option data after applying our filtering to the Berkeley Options Database are reported in Table 1.5. The average call price ranges from \$2.29 for short-term OTM to \$38.66 for long-term ITM, while the average put price is between \$2.64 for short-term OTM and \$29.43 for long-term ITM. For both calls and puts, the majority of options are ATM. Moreover, short-term options represent approximately 50% of the total sample. Table 1.5 also reports the average implied volatilities in each moneyness (S/X) and maturity category. The implied volatilities of short-term call and put options present a smile in moneyness. For medium and long-term options, implied volatility is increasing in moneyness for calls and decreasing for puts. Implied volatilities are also mostly increasing in time to expiration. The differences in implied volatilities are more pronounced for short-term options.

Option Price Bounds and Implied γ - Berkeley Options Database

We calculate price bounds from the underlying returns for each option in our sample. Table 1.6 reports, for the Berkeley Options Database, the percentage of times that observed option prices are contained in the bounds, the percentage of upper and lower bound violations and the average tightness of the bounds, for each option category. Focusing first on the aggregate results, almost all option prices (97.68% of the calls and 94.40% of the puts) are consistent with the SSD bounds. Most of the observed prices are also contained in the tighter MD bounds (95.42% of the calls and 91.71% of the puts). The larger number of violations than for the OptionMetrics database can be explained by the presence of the 1987 crash in the Berkeley data. Before the crash, there were more violations as investors were pricing options using the Black-Scholes formula, ignoring information from the underlying returns. Even so, the vast majority of option prices still lies within the MD and SSD bounds, confirming the evidence that option prices are mostly consistent with underlying returns when we consider incomplete markets.

Detailing the analysis at the level of option categories, the options that most violate the bounds are short-term ITM calls and ITM puts, where 87.82% and 73.66% of these options are contained by the MD bounds, respectively. Similarly to the OptionMetrics database, this may be due to the fact that ITM options are less liquid and may present some unreliable prices. On the other hand, 95.45% of the short-term OTM puts are consistent with the MD bounds. This contradicts the common notion that the left tail of the risk-neutral distribution is hard to be reconciled. As the time to maturity increases for medium- and long-term options, the percentage of options contained in the bounds increases. That is, the longer the maturity, the easier it is to reconcile the information from option prices and underlying returns.

We also calculate the MDNA bounds, identifying prices consistent with strictly positive risk-neutral measures. The MDNA bounds capture 86.62% of the calls and 82.77% of the puts, leaving unexplained a portion of the prices. However, a specific pattern appears when we look at the categories of options. ITM calls and OTM puts can be explained by strictly positive measures, with the much tighter MDNA bounds capturing most of the observed prices. In contrast, OTM calls and ITM puts clearly require risk-neutral measures with zeros in some states of nature to be priced, as the number of MDNA lower bound violations are considerably large. This indicates that to reconcile the right tail of the risk-neutral distribution, risk-neutral measures identified from the physical distribution of underlying returns need to decrease the probability mass in positive returns to make them compatible with option prices.

Table 1.7 reports the average implied γ over the 1987-1995 sample for each option category. On average, the implied γ of call options is decreasing in moneyness (S/X), while it is increasing for put options. That is, option prices present a smirk of the implied γ, indicating the existence of heterogeneous marginal investors in a segmented option market. OTM puts (and ITM calls) are priced by investors with positive prudence, convex marginal utilities and aversion to downside risk. On the other hand, OTM calls (and ITM puts) require large positive γ's to be priced, associated to concave marginal utilities (except for long-term options, for which the implied γ's are positive but smaller than one). The heterogeneity of marginal investors is more pronounced for short-term options, decreasing as the maturity increases.

Overall, the results for the Berkeley Options Database are qualitatively similar to those for the OptionMetrics data, providing additional robustness to our analysis.

Marginal Risk-Neutral Measures and Implied γ - Medium- and Long-Term Options

Figures 1.13 and 1.14 plot, for 1996-2019 and 1987-1995, respectively, the two-month moving average of the implied γ and the $\underline{\gamma}$ and $\overline{\gamma}$ defining the MDNA bounds, for medium- and long-term calls and puts with different moneyness. As can be seen, the plots follow closely the same patterns carefully discussed in the main paper. The only difference is that the heterogeneity of marginal minimum dispersion risk-neutral measures gets smaller as the maturity increases. Even so, the heterogeneity is time-varying and mainly driven by the implied γ of OTM calls and ITM puts.

Tables

Table 1.5: **Summary Statistics of S&P 500 Index Options (1987-1995)**

	Call Options				Put Options			
Moneyness	Short	Medium	Long	Subtotal	Short	Medium	Long	Subtotal
	\$2.29	\$6.21	\$11.61	\$6.16	\$23.25	\$26.59	\$29.43	\$25.69
[0.90, 0.97)	14.63%	15.09%	15.63%	15.06%	14.08%	14.69%	15.66%	14.63%
	{9435}	{7758}	{6621}	{23814}	{13523}	{8499}	{6735}	{28757}
	\$9.32	\$16.27	\$23.92	\$14.10	\$8.58	\$13.16	\$17.27	\$11.58
[0.97, 1.03)	14.19%	15.57%	16.31%	14.98%	13.96%	15.61%	16.48%	14.92%
	{18966}	{9086}	{7256}	{35308}	{18917}	{9165}	{7423}	{35505}
	\$27.16	\$32.48	\$38.66	\$31.00	\$2.64	\$6.35	\$10.29	\$5.34
[1.03, 1.10]	19.04%	17.67%	17.50%	18.34%	17.59%	17.58%	17.65%	17.60%
	{14494}	{7740}	{5774}	{28008}	{15357}	{8303}	{6685}	{30345}
	\$13.80	\$18.20	\$24.10	\$17.36	\$10.82	\$15.38	\$18.96	\$13.87
Subtotal	15.93%	16.08%	16.43%	16.08%	15.16%	15.94%	16.59%	15.69%
	{42895}	{24584}	{19651}	{87130}	{47797}	{25967}	{20843}	{94607}

This table presents summary statistics of the S&P 500 index option data after applying our filtering to the Berkeley Options Database. The sample ranges from January 2, 1987 to December 29, 1995. The columns Short, Medium and Long refer to the maturity categories. For each moneyness (S/X) and maturity category, the first row depicts the average option price, the second row the average implied volatility and the third row the number of observations (in braces). The average of the daily values of the S&P 500 index and the (annualized) risk-free rate in the sample period were 383.92 and 5.51%, respectively.

Table 1.6: **Option Price Bounds for S&P 500 Options (1987-1995)**

Category	MDNA Bounds In	MDNA Bounds Upper	MDNA Bounds Lower	MD Bounds In	MD Bounds Upper	MD Bounds Lower	SSD Bounds In	SSD Bounds Upper	SSD Bounds Lower
Panel A: Calls									
Short OTM	47.94%	9.58% (1.36)	42.48% (0.85)	88.09%	9.58% (1.56)	2.33% (0.37)	92.57%	5.71% (1.67)	1.72% (0.34)
Short ATM	89.41%	3.76% (1.18)	6.83% (0.93)	95.95%	3.76% (1.19)	0.28% (0.76)	98.51%	1.28% (1.27)	0.21% (0.75)
Short ITM	85.10%	9.33% (1.05)	5.57% (0.96)	87.82%	9.33% (1.05)	2.85 % (0.94)	93.86%	3.44% (1.10)	2.70% (0.94)
Medium OTM	75.81%	1.37% (1.38)	22.83% (0.88)	98.41%	1.37% (1.50)	0.22% (0.45)	99.72%	0.19% (1.70)	0.09% (0.44)
Medium ATM	99.48%	0.14% (1.18)	0.37% (0.90)	99.85%	0.14% (1.18)	0.01% (0.76)	99.97%	0.02% (1.32)	0.01% (0.75)
Medium ITM	97.62%	0.34% (1.08)	2.04% (0.94)	98.85%	0.34% (1.08)	0.81% (0.90)	99.19%	0.03% (1.19)	0.78% (0.90)
Long OTM	94.31%	0.69% (1.36)	5.00% (0.84)	99.23%	0.69% (1.39)	0.07% (0.51)	99.86%	0.06% (1.63)	0.07% (0.49)
Long ATM	99.89%	0.05% (1.19)	0.05% (0.88)	99.94%	0.06% (1.19)	0.00% (0.75)	100%	0.00% (1.38)	0.00% (0.75)
Long ITM	98.60%	0.05% (1.11)	1.35% (0.91)	99.12%	0.05% (1.11)	0.83% (0.86)	99.13%	0.04% (1.27)	0.83% (0.86)
All Calls	86.62%	3.64% (1.19)	9.74% (0.91)	95.42%	3.64% (1.23)	0.94% (0.72)	97.68%	1.50% (1.36)	0.82% (0.71)
Panel B: Puts									
Short OTM	92.07%	2.49% (1.90)	5.44% (0.70)	95.45%	2.49% (1.91)	2.06% (0.50)	97.65%	0.27% (2.66)	2.08% (0.51)
Short ATM	85.02%	3.31% (1.20)	11.67% (0.91)	92.33%	3.31% (1.20)	4.35% (0.77)	93.92%	1.85% (1.31)	4.23% (0.77)
Short ITM	40.59%	21.24% (1.02)	38.17% (0.98)	73.66%	21.24% (1.03)	5.10% (0.95)	78.95%	16.29% (1.05)	4.76% (0.95)
Medium OTM	97.10%	2.36% (1.49)	0.54% (0.68)	97.43%	2.36% (1.49)	0.20% (0.49)	99.72%	0.08% (2.16)	0.20% (0.49)
Medium ATM	96.80%	2.52% (1.20)	0.68% (0.86)	97.20%	2.52% (1.20)	0.28% (0.69)	98.67%	1.05% (1.40)	0.28% (0.69)
Medium ITM	74.54%	4.52% (1.07)	20.94% (0.96)	94.76%	4.52% (1.08)	0.72% (0.88)	96.51%	2.77% (1.13)	0.72% (0.88)
Long OTM	93.34%	6.55% (1.41)	0.10% (0.61)	93.42%	6.55% (1.41)	0.02% (0.45)	99.67%	0.30% (2.06)	0.02% (0.45)
Long ATM	94.73%	5.02% (1.24)	0.24% (0.80)	94.92%	5.02% (1.24)	0.05% (0.64)	98.01%	1.94% (1.53)	0.05% (0.64)
Long ITM	89.95%	5.69% (1.12)	4.37% (0.92)	94.03%	5.69% (1.12)	0.28% (0.81)	95.12%	4.60% (1.23)	0.28% (0.81)
All Puts	82.77%	6.22% (1.35)	11.00% (0.82)	91.71%	6.22% (1.33)	2.07% (0.69)	94.40%	3.59% (1.64)	2.01% (0.69)

This table presents empirical results for the minimum dispersion no-arbitrage (MDNA), minimum dispersion (MD) and stochastic dominance (SSD) option price bounds for S&P 500 options. For each category, column In reports the percentage of prices inside the bounds. The column Upper reports the percentage of violations of the upper bounds and, in parenthesis, the average ratio of the upper bound over the observed option price for the prices inside the bounds. Analogously, Lower reports the percentage of violations of the lower bounds and, in parenthesis, the average ratio of the lower bound over the observed option price for the prices inside the bounds. The sample ranges from January 2, 1987 to December 29, 1995.

Table 1.7: **S&P 500 Options Implied γ (1987-1995)**

	Call Options			Put Options		
Moneyness	Short	Medium	Long	Short	Medium	Long
$[0.90, 0.97)$	6.62	3.13	0.36	7.71	2.99	0.17
$[0.97, 1.03)$	0.48	-0.69	-0.83	0.76	-0.88	-1.00
$[1.03, 1.10]$	-0.61	-0.83	-0.94	-0.08	-1.15	-1.23

This table presents the average implied γ of S&P 500 index options for each moneyness (S/X) and maturity category. The sample ranges from January 2, 1987 to December 29, 1995. The columns Short, Medium and Long refer to the maturity categories.

Figures

Figure 1.13: **S&P 500 Medium- and Long-Term Options Implied γ Over Time (1996-2019)**

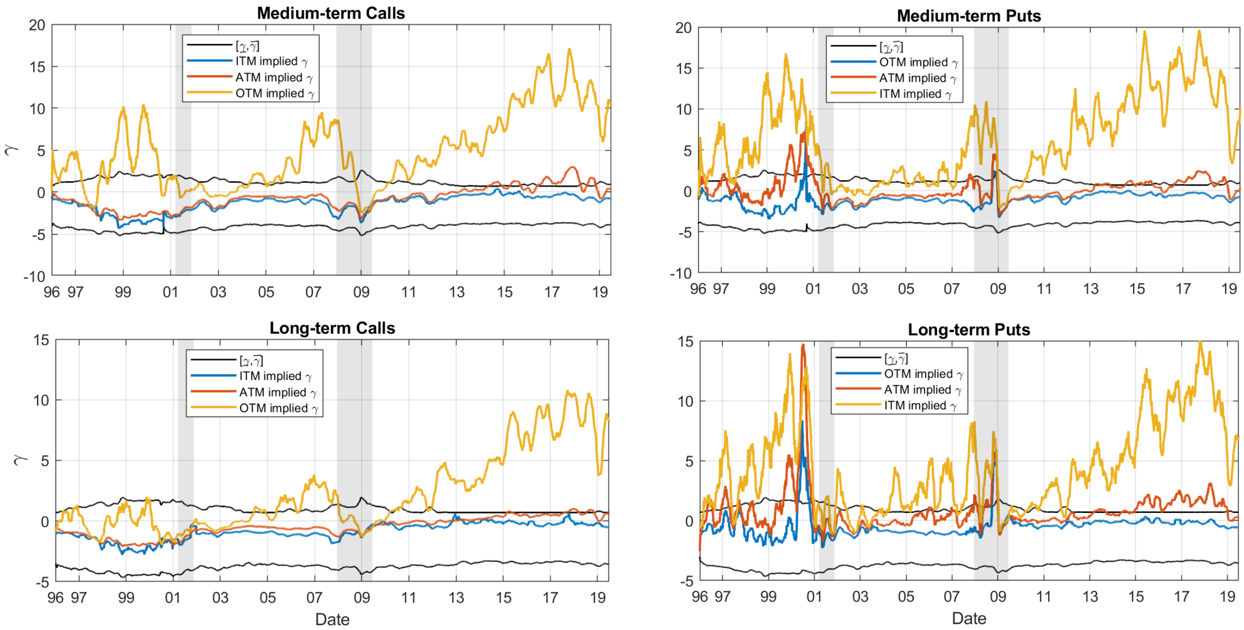

This figure plots the 2-month moving averages of the mean implied γ for OTM, ATM and ITM options and the mean $\underline{\gamma}$ and $\overline{\gamma}$ defining the MDNA bounds, for medium- and long-term calls and puts. Shaded areas depict NBER recession dates. The sample ranges from January 4, 1996 to June 28, 2019.

Figure 1.14: **S&P 500 Medium- and Long-Term Options Implied γ Over Time (1987-1995)**

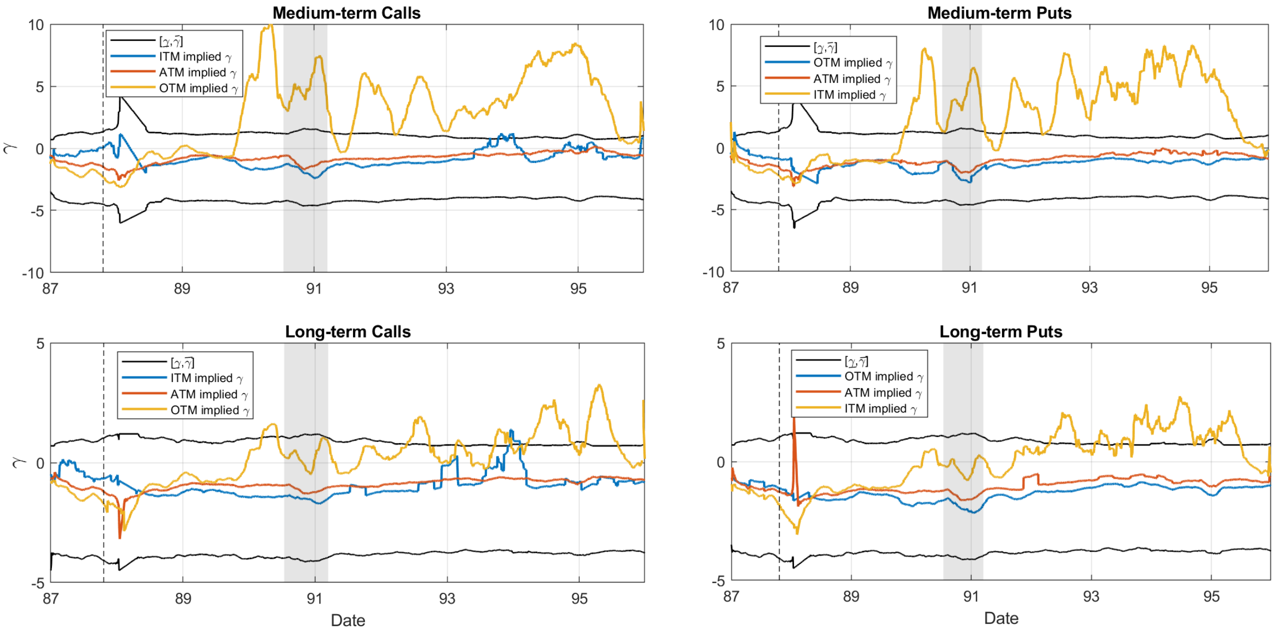

This figure plots the 2-month moving averages of the mean implied γ for OTM, ATM and ITM options and the mean $\underline{\gamma}$ and $\overline{\gamma}$ defining the MDNA bounds, for medium- and long-term calls and puts. Shaded areas depict NBER recession dates and the vertical dashed line corresponds to the October 1987 market crash. The sample ranges from January 2, 1987 to December 29, 1995.

Chapter 2

Nonparametric Option Pricing with Generalized Entropic Estimators

When pricing options using only information from underlying returns, one naturally faces incomplete markets where there exists a multiplicity of risk-neutral measures. We nonparametrically estimate a family of risk-neutral measures that are the “closest” possible to the physical distribution, in the sense of minimizing generalized entropy. Each measure distorts higher moments of the physical probabilities in a particular way. With Monte Carlo experiments, we study the option pricing implications of this family under different degrees of market incompleteness. In a large-scale empirical application, we compare the out-of-sample option pricing performance of the generalized entropic estimators with benchmark methods. We identify the estimators in the family that outperform the Black-Scholes and GARCH option models for different options in the cross-section.

This chapter is co-authored with Caio Almeida, Kym Ardison and Rafael Azevedo.

2.1 Introduction

The seminal Black and Scholes (1973) and Merton (1973) model allows for a closed-form solution to the price of a European option. This is possible due to strong parametric assumptions that are unlikely to hold in practice. In fact, several predictions of the model have been rejected by data on the S&P 500 option market, which comes closest to satisfying the model assumptions. For instance, implied volatilities of S&P 500 options are not constant across strike prices and maturities (Rubinstein, 1994; Derman and Kani, 1994), and the risk-neutral distribution of the market index is more skewed to the left and has fatter tails than the lognormal distribution (Jackwerth and Rubinstein, 1996; Ait-Sahalia and Lo, 1998).

One of the shortcomings of the Black-Scholes model is the constant volatility assumption. Derman and Kani (1994), Dupire (1994) and Rubinstein (1994) consider the possibility that local volatility is a deterministic function of the asset price and time. This function is estimated as a binomial lattice that exactly fits a cross-section of observed option prices. A different strand of the literature develops more general continuous-time parametric models, accounting for additional sources of risk such as stochastic volatility (Heston, 1993) and jumps (Bates, 2000). These models are difficult to implement and their parameters are usually estimated using option prices (Bakshi, Cao and Chen, 1997; Broadie, Chernov and Johannes, 2007). Other papers adopt discrete-time GARCH option models, where the parameters governing the time-varying volatility and the option pricing formula can be computed directly from the historical underlying returns (Duan, 1995; Heston and Nandi, 2000; Christoffersen, Heston and Jacobs, 2013). Observed option prices can also be used in the estimation procedure.

Rather than relying on parametric assumptions, alternative methods have been proposed to nonparametrically estimate option prices. Examples include neural networks (Hutchinson, Lo and Poggio, 1994; Garcia and Gençay, 2000), kernel regression (Ait-Sahalia and Lo, 1998) and locally polinomial regression (Ait-Sahalia and Duarte, 2003). These techniques make use of large amounts of option market data to curve-fit the true option pricing function. In a different approach, Stutzer (1996) nonparametrically estimates the risk-neutral measure that is "closest" to the physical distribution, in the sense of minimizing entropy, while correctly pricing underlying returns.[1] Option prices can then be calculated by their expected discounted payoffs under the risk-neutral distribution. Unlike the aforementioned procedures, this approach is not a curve-fitting technique, but rather constitutes a predictive theory of pricing that formalizes the risk-neutralization process. Therefore, it does not strictly require the use of option data.

While there has been indisputable success for option pricing techniques leverag-

[1]It is also possible to force the risk-neutral probabilities to be consistent with observed option prices.

ing the information from rich cross-sections of options, there are important unresolved challenges when pricing options using only underlying returns. In such cases, one naturally faces incomplete markets where there exists an infinity of admissible risk-neutral measures.[2] The incompleteness begs the question of how to choose the appropriate risk-neutral measure. This is important not only for option pricing in illiquid option markets, but also to shed light on the degree of incompleteness of an option market with respect to the underlying asset.

In this paper, we contribute to fill this gap. We study the option pricing implications of a comprehensive family of risk-neutral measures. Each measure minimizes a different Cressie and Read (1984) loss function measuring the generalized entropy with respect to the physical distribution, while correctly pricing the underlying returns.[3] This procedure is based on the premise that there is no *a priori* reason for a risk-neutral measure to deviate from the physical distribution other than satisfying the martingale property. The Cressie-Read family embeds as particular cases several divergence measures, including entropy. Therefore, our method can be seen as a generalization of the predictive theory of pricing in Stutzer (1996). In particular, this generalization recognizes the potential incompleteness of the market by considering a family of generalized entropic estimators.

We start the analysis by investigating the option pricing implications of the generalized entropic estimators with Monte Carlo experiments in simulated economies. This allows us to understand how the estimators behave under different sources of risk and market incompleteness. The first economy is a Black and Scholes (1973) environment, where there is only one source of risk given by the underlying price process. In this case, we show that there is an optimal Cressie-Read risk-neutral measure that matches the risk-neutral distribution of the Black-Scholes model. The optimal estimator is completely determined by the model parameters governing risk premium. The second economy additionally incorporates stochastic volatility and jumps (Bates, 2000; Duffie, Pan and Singleton, 2000). In this environment, the optimal generalized entropic estimator for a given option depends not only on the levels of risk premia in the economy, but also on the option moneyness and maturity. This highlights the importance of considering a family of estimators to price options in incomplete markets.

In a large-scale empirical application with 1,817,095 options on the S&P 500 index from January 4, 1996 to June 28, 2019, we investigate the out-of-sample pricing performance of the generalized entropic estimators. We consider as benchmarks the Black-Scholes model and GARCH option models, which can be estimated using only information

[2] i.e., measures that correctly price the underlying asset but potentially give different prices to derivatives not spanned by the underlying.

[3] Kitamura (1996), Baggerly (1998) and Newey and Smith (2004) suggest using this comprehensive family to measure the divergence between distributions.

from underlying returns. We find that it is always possible to improve on the performance of the benchmark methods by considering distinct generalized entropic estimators for different options in the cross-section. We identify these estimators, where the improvement in performance mainly comes from different compensations for the tails of the risk-neutral distribution.

The remainder of the paper is organized as follows. Section 2.2 derives the non-parametric option pricing methodology under a generalized entropic hypothesis. Section 2.3 reports the results for the Monte Carlo experiments. Section 2.4 discusses the findings of the empirical application pricing S&P 500 options out-of-sample. Section 2.5 concludes the paper.

2.2 Option Pricing under a Generalized Entropic Hypothesis

Given a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, consider pricing derivatives on a certain underlying asset, whose prices are observed under the physical measure $\mathbb{P}$. The assumption of absence of arbitrage guarantees the existence of at least one risk-neutral measure $\mathbb{Q}$ equivalent to $\mathbb{P}$ under which the discounted price of any asset is a martingale (see Duffie, 2001). In particular, considering a European call option expiring in T days, its price at time t is given by:

$$C = \mathbb{E}_t^{\mathbb{Q}} \left[\frac{\max\left(p_t R - X, 0\right)}{(1 + r_f)^T} \right], \tag{2.1}$$

where $\mathbb{E}_t^{\mathbb{Q}}$ is the conditional expectation taken with respect to a risk-neutral measure $\mathbb{Q}$, X is the option strike, r_f is the daily risk-free rate, p_t is the current price of the underlying asset, R is the T-day return from the physical distribution and $p_t R \equiv p_T$ is the price of the underlying at time T. The sample counterpart to (2.1) is given by an expectation over $k = 1, ..., n$ states of nature, for which we have returns R_k drawn from the physical distribution with weights $\pi_k = 1/n$:

$$C = \sum_{k=1}^{n} \pi_k^{\mathbb{Q}} \left[\frac{\max\left(p_t R_k - X, 0\right)}{(1 + r_f)^T} \right], \tag{2.2}$$

where $\pi_k^{\mathbb{Q}}$ is the risk-neutral counterpart of the empirical measure π_k and must correctly price a set of basis assets, or, equivalently, satisfy the martingale property. In general, the number of states of nature n is larger than the number of basis assets, so the market is incomplete and there exists multiple risk-neutral measures under no-arbitrage. The pricing problem then becomes how to properly choose one specific measure $\pi_k^{\mathbb{Q}}$ from the set of existing risk-neutral measures.

Stutzer (1996) suggests to choose the risk-neutral measure that minimizes the entropy, or the Kullback-Leibler information criterion (KLIC), between the risk-neutral and physical distributions. This procedure is based on the idea that it is reasonable to select a risk-neutral measure that incorporates no other information than the one available from the physical distribution and the restrictions defining the martingale property. Pricing options under the minimum entropy risk-neutral measure constitutes a predictive theory of pricing based on the entropic hypothesis, i.e., the use of the physical distribution as a prior to a minimum KLIC posterior risk-neutral distribution.

In this paper, we adopt the same premise that there is no *a priori* reason for a risk-neutral measure to deviate from the physical distribution other than to satisfy the martingale property. This naturally suggests to choose risk-neutral measures that are the "closest" possible to the physical probabilities. The sense of closeness comes from minimizing a measure of divergence between the two probability measures. However, the divergence can be measured by any convex loss function. We derive a family of nonparametric estimators for an option price by considering risk-neutral measures minimizing a comprehensive family of discrepancy loss functions (Cressie and Read, 1984). In other words, we propose a predictive theory of option pricing under a generalized entropic hypothesis: the physical measure represents a prior to a family of posterior risk-neutral measures minimizing different discrepancies. This is consistent with the multiplicity of admissible measures in an incomplete market.

The loss functions in the Cressie-Read family measuring the discrepancy between $\mathbb{Q}$ and $\mathbb{P}$ are defined by:

$$\phi_\gamma\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right) = \frac{\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)^{\gamma+1} - 1}{\gamma(\gamma+1)}, \quad \gamma \in \mathbb{R}. \tag{2.3}$$

Kitamura (1996), Baggerly (1998) and Newey and Smith (2004) suggest using this family, which includes as particular cases several well-known divergence measures. For instance, the Euclidean divergence ($\gamma = 1$), the KLIC ($\gamma \to 0$), the Hellinger divergence ($\gamma = -1/2$), the empirical likelihood ($\gamma \to -1$) and Pearson's Chi-Square ($\gamma = -2$). As shown in Chapter 1, each discrepancy takes into account different sensitivities to higher moments of the risk-neutral probabilities. This implies that if skewness, kurtosis and tail probabilities are relevant for option pricing, the generalized entropic estimators will likely be able to capture these higher moments.

Let $\mathbf{R}$ denote a K-dimensional random vector on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ representing the gross returns of a set of K basis assets. These assets can be, for instance, the underlying and options with observed prices, or just the underlying asset. Suppose further that there is a risk-free rate R_f and let $\mathbf{R}_f = R_f \mathbf{1}_K$, where $\mathbf{1}_K$ is a K-vector of ones. The generalized entropic risk-neutral measures solve the following minimization in the space of admissible measures with $I^{\phi_\gamma}(\mathbb{Q},\mathbb{P}) \equiv \mathbb{E}\left[\phi_\gamma(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}})\right] < \infty$:

$$\mathbb{Q}^*_\gamma = \underset{Q}{arg\ min}\ \mathbb{E}\left[\phi_\gamma\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)\right] \equiv \int \phi_\gamma\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)\mathrm{d}\mathbb{P},\ \text{ s.t. }\ \mathbb{E}^{\mathbb{Q}}\left[\mathbf{R}-\mathbf{R}_f\right] = \mathbf{0}, \tag{2.4}$$

where, under the restriction that $\mathbb{Q}$ is absolutely continuous with respect to $\mathbb{P}$, $\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}$ is a Radon-Nikodym derivative[4]. The risk-neutral measure must also be nonnegative and integrate to one. By definition, we have that any measure $\mathbb{Q}$ satisfies $I^{\phi_\gamma}(\mathbb{Q},\mathbb{P}) \geq I^{\phi_\gamma}(\mathbb{Q}^*_\gamma,\mathbb{P})$. Furthermore, the physical measure $\mathbb{P}$ represents a prior to the risk-neutral measure $\mathbb{Q}$, while the restrictions in (2.4) define the martingale property. With data on basis assets returns available, these constraints can be imposed and used to update the prior according to a measure of information gain determined by ϕ_γ. That is, we are looking for the risk-neutral measure that adds the minimum amount of information, according to ϕ_γ, needed for $\mathbb{P}$ to satisfy the martingale property. Note also that $I^{\phi_\gamma}(\mathbb{Q},\mathbb{P}) \geq 0$, with equality only when $\mathbb{Q} = \mathbb{P}$. Therefore, we have $I^{\phi_\gamma}(\mathbb{Q}^*_\gamma,\mathbb{P}) = 0$ only in the case that $\mathbb{P}$ already satisfies the constraints in (2.4).

At first glance, the variational problem (2.4) might seem difficult to solve. However, Chapter 1 shows, using the results in Almeida and Garcia (2017), that the minimum discrepancy problem can be solved through a much simpler finite dimensional dual problem. In order to empirically obtain the risk-neutral measures minimizing the Cressie-Read family of discrepancies, we consider the sample version of problem (2.4). The sample space Ω is discrete and finite, with states of nature $k = \{1, ..., n\}$, where $n > K$. Let $\{\mathbf{R}_k\}_{k=1}^n$ be the observed gross returns of the K basis assets, where each $\mathbf{R}_k$ is independent and identically distributed according to $\mathbb{P}$. The unknown physical measure $\mathbb{P}$ can be replaced by the empirical measure $\mathbb{P}_n$ that gives weights $\pi_k = 1/n$ to the realization of each state of nature.[5] This allows us to exchange the expectation $\mathbb{E}$ with its sample counterpart $\frac{1}{n}\sum_{k=1}^n \equiv \sum_{k=1}^n \pi_k$. In the following corollary from Chapter 1, we summarize the results for the sample version of the problem of finding a minimum dispersion Cressie-Read risk-neutral measure:

Corollary 5. *Consider the primal problem:*

$$\begin{aligned} &\underset{\{\pi_1^Q,\dots,\pi_n^Q\}}{min}\ \ \sum_{k=1}^n \pi_k \frac{\left(\pi_k^{\mathbb{Q}}/\pi_k\right)^{\gamma+1}-1}{\gamma(\gamma+1)}, \\ &s.t.\ \ \sum_{k=1}^n \pi_k^{\mathbb{Q}}\left(\boldsymbol{R}_k - \boldsymbol{R}_f\right) = 0,\ \ \sum_{k=1}^n \pi_k^{\mathbb{Q}} = 1,\ \pi_k^{\mathbb{Q}} \geq 0\ \ \forall\, k. \end{aligned} \tag{2.5}$$

[4] i.e., $\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} : \Omega \to [0,\infty)$ is a measurable function such that for any measurable set $A \subseteq \Omega$, $\mathbb{Q}(A) = \int_A \frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\mathrm{d}\mathbb{P}$.

[5] This constitutes an optimal nonparametric estimator for $\mathbb{P}$. For more details, see Kitamura (2006).

Absence of arbitrage in the observed sample implies that the values of the primal problem coincide (with dual attainment) with the values of the dual problems below:

i) if $\gamma > 0$:

$$\lambda_\gamma^* = \underset{\lambda \in R^K}{arg\ max} - \frac{1}{\gamma+1}\sum_{k=1}^{n} \pi_k \left(1 + \gamma\lambda' \left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)\right)^{\left(\frac{\gamma+1}{\gamma}\right)} I_{\Lambda_\gamma(\boldsymbol{R}_k)}(\lambda), \tag{2.6}$$

ii) if $\gamma < 0$:

$$\lambda_\gamma^* = \underset{\lambda \in \Lambda_\gamma}{arg\ max} - \frac{1}{\gamma+1}\sum_{k=1}^{n} \pi_k \left(1 + \gamma\lambda' \left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)\right)^{\left(\frac{\gamma+1}{\gamma}\right)}, \tag{2.7}$$

iii) if $\gamma = 0$, the maximization is unconstrained:

$$\lambda_0^* = \underset{\lambda \in R^K}{arg\ max}\ - \sum_{k=1}^{n} \pi_k\, e^{\lambda'\left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)}, \tag{2.8}$$

where $\Lambda_\gamma = \{\lambda \in \mathbb{R}^K : for\ k = 1, ..., n,\ (1 + \gamma\lambda' (\boldsymbol{R}_k - \boldsymbol{R}_f)) > 0\}$ and $\Lambda_\gamma(\boldsymbol{R}_k) = \{\lambda \in \mathbb{R}^K : (1 + \gamma\lambda' (\boldsymbol{R}_k - \boldsymbol{R}_f)) > 0\}$.

The minimum dispersion risk-neutral measure can then be recovered from the first-order conditions of (2.6), (2.7) and (2.8) with respect to λ, evaluated at λ_γ^*:

$$\pi_k^{\mathbb{Q}_\gamma^*}(\gamma, \boldsymbol{R}) = \frac{(1 + \gamma\lambda_\gamma^{*\prime}(\boldsymbol{R}_k - \boldsymbol{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\boldsymbol{R}_k)}(\lambda_\gamma^*)}{\sum_{i=1}^{n}(1 + \gamma\lambda_\gamma^{*\prime}(\boldsymbol{R}_i - \boldsymbol{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\boldsymbol{R}_i)}(\lambda_\gamma^*)},\ k = 1, ..., n;\ \gamma > 0, \tag{2.9}$$

$$\pi_k^{\mathbb{Q}_\gamma^*}(\gamma, \boldsymbol{R}) = \frac{\left(1 + \gamma\lambda_\gamma^{*\prime} \left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)\right)^{\frac{1}{\gamma}}}{\sum_{i=1}^{n} \left(1 + \gamma\lambda_\gamma^{*\prime} \left(\boldsymbol{R}_i - \boldsymbol{R}_f\right)\right)^{\frac{1}{\gamma}}}, k = 1, ..., n;\ \gamma < 0, \tag{2.10}$$

$$\pi_k^{\mathbb{Q}_0^*}(0, \boldsymbol{R}) = \frac{e^{\lambda_0^*\left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)}}{\sum_{i=1}^{n} e^{\lambda_0^*\left(\boldsymbol{R}_i - \boldsymbol{R}_f\right)}}, k = 1, ..., n;\ \gamma = 0. \tag{2.11}$$

Proof. See Chapter 1. □

By solving the dual problem specified above for each discrepancy γ, we obtain a direct set of estimates λ_γ^* for λ, each one leading to a different minimum dispersion risk-neutral measure. In the context of option pricing, the returns for the estimation of the risk-neutral measures will be compounded according to the option maturity. For the risk-free rate, this implies that $R_f = (1 + r_f)^T$. By substituting the minimum dispersion risk-neutral measures in the option pricing equation (2.2), we obtain a family of generalized entropic estimators for the option price:

$$C_\gamma = \sum_{k=1}^{n} \pi_k^{\mathbb{Q}_\gamma^*} \left[\frac{\max\left(p_t R_k - X, 0\right)}{(1 + r_f)^T}\right],\ \ \gamma \in \mathbb{R}, \tag{2.12}$$

where the parameter γ indexes the minimum dispersion risk-neutral measure and the corresponding estimator. Each measure is a nonlinear hyperbolic function of an optimal linear combination of the basis assets. For $\gamma \leq 0$, the minimum dispersion measures are strictly positive, while for $\gamma > 0$ there can be zeros in some states of nature due to the indicator function in (2.9). The measures are convex for $\gamma < 1$, concave for $\gamma > 1$ and linear for the case that $\gamma = 1$. Chapter 1 shows that, for the case of estimating the risk-neutral measure using only underlying returns, the implied option prices are monotonically decreasing in γ. This is because the smaller the γ, the more convex is the measure, giving higher weights to states of nature with extreme negative and positive underlying returns, where options pay off.

The methodology above allows for the inclusion of options among the basis assets. In that case, this approach would be similar to existing methods that curve-fit the true option pricing function from option data. Instead, we focus on investigating the option pricing implications of the family of generalized entropic estimators using only underlying returns. This can be particularly useful for option pricing in illiquid option markets, besides bringing new insights onto the degree of incompleteness of option markets.

2.2.1 Interpretation of the Generalized Entropic Estimation

Obtaining a minimum discrepancy risk-neutral measure according to a convex function ϕ can be seen as a generalized minimum contrast (GMC) estimation procedure (Bickel et al., 1993). This procedure treats the data distribution nonparametrically by comparing the empirical distribution of the data with the family of distributions implied by a statistical model. In our case, the statistical model is the set of all probability measures that satisfy the moment restrictions characterizing the martingale property for the basis assets.

There is a close relationship between GMC estimators and the class of generalized empirical likelihood (GEL) estimators (Smith, 1997). Newey and Smith (2004) show that the GEL objective function is given by the dual problem of the minimum discrepancy sample problem with Cressie-Read discrepancies. The parameter γ indexes the particular estimator in the GEL class. For instance, $\gamma \to -1$ yields the empirical likelihood (Owen, 1988), $\gamma \to 0$ the exponential tilting (Kitamura and Stutzer, 1997) and $\gamma = 1$ the continuous updating estimator (Hansen, Heaton and Yaron, 1996). These estimators are robust againts distributional assumptions, possess desirable properties analogous to those of parametric likelihood procedures and have been used to improve on the small sample properties of GMM estimators.[6]

The minimum dispersion risk-neutral measures are the implied probabilities com-

[6] See Kitamura (2006) for a review.

ing from the GMC estimation. These probabilities are useful in a number of applications. For instance, Brown and Newey (1998) conduct efficient estimation of moment conditions based on implied probabilities. Brown and Newey (2002) suggest their use on the estimation of probability distribution functions via bootstrapping schemes. Smith (2004) use implied probabilities to obtain efficient moment estimation for GEL estimators. Antoine, Bonnal and Renault (2007) show, in the context of euclidean likelihood, how these probabilities contain important information coming from overidentifying restrictions that can decrease the variance of the estimator. Almeida and Garcia (2012) make use of the implied probabilities to compare misspecified asset pricing models, while Almeida and Garcia (2017) apply them to derive nonparametric bounds for pricing kernels. We employ the Cressie-Read family of implied probabilities to price options nonparametrically using information from the underlying returns.

2.3 Monte Carlo Analysis

We start the analysis of the option pricing implications of the generalized entropic estimators by conducting Monte Carlo experiments in simulated economies coming from well-known parametric models. This allows us to investigate how the estimators behave in controlled environments with different sources of risk and market incompleteness. The parametric models considered allow for analytic solutions to the option pricing equation (2.1) that will be used as a benchmark for our nonparametric estimators. We estimate the risk-neutral measures using underlying returns coming from the model-implied physical distribution, and compare them with the theoretical option prices implied by risk-neutral parameters of the model. We will consider pricing European call options for several combinations of moneyness and maturity. The results for put options are analogous given put-call parity.

The first economy considered is a Black and Scholes (1973) environment, while the second economy comes from the stochastic volatility and correlated jumps (SVCJ) model (Bates, 2000; Duffie, Pan and Singleton, 2000). The SVCJ model captures important empirical stylized facts in equity markets and reliable estimates of its physical and risk-neutral parameters have been obtained in the option pricing literature. Appendix B 1.8 in Chapter 1 presents in detail the models, the simulation procedure and the parameters for the Black-Scholes model. Table 2.1 reports the parametrization for the SVCJ model.

We analyze two statistics based on the option pricing error probability distribution: the mean percentage pricing error (MPE) and the mean absolute percentage pricing error (MAPE).[7] The MPE approximates the expected relative difference between the es-

[7] The MPE is given by $\sum_{j=1}^{J} \frac{C_{\gamma,j} - C_{\text{model}}}{C_{\text{model}}}$, while the MAPE is given by $\sum_{j=1}^{J} \frac{|C_{\gamma,j} - C_{\text{model}}|}{C_{\text{model}}}$, where

timated and the true option price. The MAPE reflects the accuracy of the option pricing procedure. For a given parametric model, we draw 200 underlying returns, compounded according to the option maturity, from the model-implied physical distribution. From the returns, we estimate the minimum dispersion risk-neutral measures for different values of γ and price options according to equation (2.12). This procedure is repeated 5000 times in order to calculate the MPE and MAPE. The pricing errors are calculated with respect to the theoretical option price implied by the risk-neutral parameters of the model.

2.3.1 Black-Scholes Environment

We begin with a theoretical result showing that in a Black-Scholes environment there is an optimal generalized entropic estimator completely determined by the parameters governing risk premium in the model.

Proposition 11. *Suppose that the underlying asset dynamics comes from the Black-Scholes model with drift μ, risk-free rate r and volatility σ. Then, there is an optimal Cressie-Read discrepancy indexed by γ^* for which the implied minimum dispersion risk-neutral measure coincides with the risk-neutral measure from the Black-Scholes model, given by:*

$$\gamma^* = -\frac{\sigma^2}{\mu - r}. \tag{2.13}$$

Proof. See Appendix A 2.6. □

The proposition above indicates that, in a Black-Scholes environment, the equity risk premium determines the optimal estimator for any option, regardless of the moneyness and maturity. Smaller equity premia and higher volatilities imply that options require more convex risk-neutral measures giving more weight to extreme outcomes of underlying returns. Moreover, if the equity premium goes to zero, i.e., $\mu - r \to 0$, the optimal discrepancy is given by $\gamma^* \to -\infty$. This is consistent with the fact that when there is no risk premia in the economy, the physical measure already satisfies risk neutrality, implying that the minimum dispersion risk-neutral measure is precisely $\mathbb{P}$. Conversely, an arbitrarily large equity premium or small volatility implies $\gamma^* \to 0$. Furthermore, the proposition shows that, in a Black-Scholes economy, generalized entropic estimators with $\gamma > 0$ should not be considered if $\mu > r$.

We now turn to the simulation results. Note that, for the model parameters $\mu = 0.1$, $r = 0.05$ and $\sigma = 0.2$, we obtain $\gamma^* = -0.8$. Figure 2.1 depicts the MPE for a range of γ's and call options with different combinations of moneyness and maturity. As can be seen, all MPE curves cross the zero line around $\gamma = -0.8$, corroborating the

C_{model} is the true option price given by the model and $j = 1, ..., J$ indexes each simulation.

theoretical result.[8] Given that implied option prices are monotonically decreasing in γ, the MPE also decreases monotonically with γ. This sheds light on previous findings in the literature. Stutzer (1996) and Grey and Newman (2005) price options with the risk-neutral measure that minimizes the KLIC, while Haley and Walker (2010) also consider two other estimators in the Cressie-Read family: the Euclidean divergence ($\gamma = 1$) and empirical likelihood ($\gamma = -1$). In the same Black-Scholes environment as considered here, they document a consistent large negative bias for the estimated option price by Euclidean divergence, followed by a smaller negative bias for the KLIC measure and a small positive bias for the empirical likelihood estimator. Proposition 11 and Figure 2.1 reveal that these results are not general properties of the estimators as previously believed, but rather an implication of the Black-Scholes model parametrization considered.

Figure 2.1 also shows that, for a given option, the relative difference in performance across estimators, or how steep is the MPE curve, depends on the moneyness and maturity. In particular, the difference is increasing in the time to maturity and decreasing in moneyness (S/X). This is depicted in more detail in the MPE surfaces in Figure 2.2. For OTM calls, the MPE varies considerably with γ, while for ITM calls the MPE varies less across estimators. For a fixed γ, the estimator performance is better for ITM calls. As for the accuracy of the generalized entropic estimators, Figure 2.3 reports the MAPE surfaces for a range of γ's and different options. The MAPE is decreasing in moneyness. OTM calls are associated with higher MAPE because only a few large positive returns make these options pay at expiration, and the variation in the occurrence of such returns across simulations has a large impact on the estimated option prices. MAPE also tends to be higher the less time to maturity the option has. For one-month options, there is little difference in MAPE across γ's. For options with time to maturity greater than or equal to 3 months, the MAPE is U-shaped with respect to γ, where γ's close to -0.8 give the smallest MAPE.

In sum, in a Black-Scholes environment, where there is only one source of risk coming from the price process for the underlying asset, we show how to obtain an optimal generalized entropic estimator that correctly prices in population any option using only information from the underlying returns. In particular, the optimal estimator depends on the parameters governing risk premium in the Black-Scholes model.

2.3.2 SVCJ Environment

The SVCJ model presents additional sources of risk that account for stylized facts in real markets: stochastic volatility and jumps. In this environment, we start by investigating the relation between the generalized entropic estimators and theoretical option

[8] We also tested for other parametrizations of the Black-Scholes model and this result always holds.

prices associated to different parametrizations of risk premium in the additional risk sources. While our estimators use underlying returns from the physical distribution, we calculate the theoretical option prices according to four different parametrizations of the risk-neutral measure as estimated by Broadie, Chernov and Johannes (2007), accounting for different combinations of risk premia. More specifically, they consider four cases allowing for the existence or not of volatility risk premium and volatility of price jumps risk premium, and estimate the risk-neutral parameters for each case. For instance, the $\mathbb{Q}_2$ parametrization in Table 2.1 allows for both, i.e., $\kappa_v^{\mathbb{Q}} \neq \kappa_v$ and $\sigma_s^{\mathbb{Q}} \neq \sigma_s$.

Figure 2.4 reports the MPE for a range of γ's and different call options for each risk-neutral parametrization. Again, MPE is monotonically decreasing in γ and, in general, there is an estimator that gives zero MPE for each option. In contrast to the Black-Scholes environment, however, the optimal estimator depends not only on the risk premia of the economy, but also on the option moneyness and maturity. This indicates that in the presence of sources of market incompleteness such as stochastic volatility and jumps, distinct generalized entropic estimators will be needed to price different options in the cross-section. Even so, for most options the optimal estimator is in the γ region between -4 and -1. The exceptions are mainly one-month options, which require more negative γ's to be priced. As for the relative difference in performance across estimators, we can see that it is increasing in the time to maturity and decreasing in moneyness (S/X). On the other hand, the difference between the MPE curves of different risk premia tends to be smaller for longer maturities.

In order to assess the accuracy of the generalized entropic estimators, Figure 2.5 plots the MAPE surfaces for a range of γ's and different options, considering the model risk-neutral parametrization that allows for all sources of risk premia. The patterns are similar to the Black-Scholes environment. MAPE tends to be decreasing in moneyness and time to maturity. For one-month options, there is little difference in MAPE across γ's. For options with time to maturity greater than or equal to 3 months, the MAPE is U-shaped with respect to γ, where γ's close to -1 give the smallest MAPE.

Overall, the results above indicate that in the presence of additional sources of risk and market incompleteness such as stochastic volatility and jumps, the optimal generalized entropic estimator for a given option will depend not only on the levels of risk premia in the economy, but also on the option moneyness and maturity. This suggests that we should consider distinct estimators in the family to price different options in the cross-section.

2.4 Empirical Application

The results from the previous section suggest that, in the presence of sources of market incompleteness, we should consider distinct generalized entropic estimators to price different options in the cross-section. In this section, we investigate the out-of-sample pricing performance of the following estimators: Pearson's Chi-Square ($\gamma = -2$), empirical likelihood ($\gamma = -1$), Hellinger ($\gamma = -0.5$), KLIC ($\gamma = 0$), Euclidean divergence ($\gamma = 1$) and $\gamma = 5$. In particular, we price options on the S&P 500 index using only returns on the underlying.[9] We also consider as a benchmark other option pricing methods that can be estimated using only information from underlying returns.

2.4.1 Data and Estimation

The data is obtained from OptionMetrics, for the sample period ranging from January 4, 1996 to June 28, 2019. The data consists on bid and ask quotes, volume, strike, expiration date and open interest for each option. Daily data on the S&P 500 index is obtained from Bloomberg, for the sample covering from July 1, 1954 to June 28, 2019. We consider the 3-month Treasury Bill rate from the Saint Louis FRED database as a proxy for the risk-free rate, for the same period for which the option data is available. As standard in the literature, we use the mid-point of the bid and ask quotes as the option prices and apply a handful of filters to the raw data. Observations with zero volume, with bid-ask spread lower than 0.85, with price lower than 1/8 and implied volatility greater than 0.7 are dropped. We also eliminate options that violate the no-arbitrage bounds of Merton (1973). To calculate the time to expiration of each option, we take into account if the contract settlement is at the open or close.[10] Since the S&P 500 index typically pays a dividend, we estimate the dividend yield from the put-call parity relation, using the pair of call and put options that are closest to ATM.[11]

We focus on options with time to expiration between 20 and 365 calendar days and moneyness between 0.90 and 1.10. The option data is divided in six moneyness-maturity categories. With respect to moneyness, we group options in three intervals: OTM call (ITM put) if $S/X \in [0.90, 0.97)$, ATM if $S/X \in [0.97, 1.03)$ and ITM call (OTM put) if

[9] S&P 500 options are one of the most traded derivatives in the world, the most actively traded European options and have been the focus of many previous studies, such as Rubinstein (1994), Jackwerth and Rubinstein (1996), Bakshi, Cao and Chen (1997), Ait-Sahalia and Lo (1998), Garcia and Gençay (2000), Heston and Nandi (2000) and Christoffersen, Heston and Jacobs (2013).

[10] For PM settled options, the time to expiration is the number of days between the trade date and the expiration date. For AM settled options, we use the number of days between the dates less one.

[11] For every day in the sample and each maturity T, we identify the pair of ATM call and put option prices, c and p, and estimate the dividend yield as $q = -(1/T)\ln[(c - p + Xe^{-rT})/S]$, where r is the corresponding continuously compounded risk-free rate. This procedure is usually preferred to the backward-looking approach of estimating the future dividend rate using past daily dividend payments on the index.

$S/X \in [1.03, 1.10]$. As for maturities, we classify options as short-term ($[20, 90)$ days) and long-term ($[90, 365]$ days). Our final sample consists of 1,817,095 options, with 818,666 calls and 998,429 puts. There is an average of 307 options in the cross-section per trading day. Summary statistics are reported in Table 2.2. The average call price ranges from \$8.19 for short-term OTM to \$148.67 for long-term ITM, while the average put price is between \$13.49 for short-term OTM and \$133.19 for long-term ITM. Most calls are ATM, while most puts are OTM.

Table 2.2 also reports the average implied volatilities in each moneyness-maturity category. The implied volatilities of call options are increasing in both moneyness and time to expiration. Put implied volatilities are also increasing in time to expiration, while decreasing in moneyness for long-term options. Short-term puts present a U-shaped pattern as the option goes from OTM to ATM and then to ITM, where the OTM put implied volatilities take the highest values. The implied volatility biases are more pronounced for short-term options. This indicates that such options are more mispriced by the Black-Scholes model and are potentially harder to price.

We conduct four empirical exercises comparing the performance of the generalized entropic estimators to benchmark methods. We price options each day in our sample, where all models are estimated using only information up to the previous day. Therefore, all results are out-of-sample. The first benchmark is the Black-Scholes model with the volatility estimated as the annualized standard deviation of the underlying returns in the past 90 days. To compare our method with this approach, we estimate the unconditional underlying return distribution as the histogram of past overlapping T-day returns calculated in the index sample (where T is the time to maturity of the option), and make it conditional by adjusting its volatility to the standard deviation of the returns in the past 90 days. Using the resulting underlying returns, we estimate the generalized entropic risk-neutral measures and price each option according to equation (2.12).

The second benchmark is the Black-Scholes model with the volatility estimated as the implied volatility of the ATM option in the previous day.[12] To compare our method with this approach, we again estimate the unconditional underlying return distribution as the histogram of past overlapping T-day returns calculated in the index sample. We make this distribution conditional by adjusting its volatility to the previous day ATM implied volatility and discounting a premium to obtain the physical volatility, following footnote 37 of Chapter 1. We use the resulting underlying returns in the generalized entropic estimators.

Both exercises above compare the generalized entropic estimators with the Black-Scholes model, which is known to be misspecified. An alternative benchmark that can be

[12] In the previous day, we interpolate the ATM implied volatilities over the different available maturities and use the interpolated value for the maturity of each option in the current day.

easily estimated using only information from underlying returns is the class of GARCH option models. These models start by specifying the physical process for the underlying returns as a discrete-time GARCH process with time-varying volatility. Assuming an exponential change of measure, it is then possible to obtain analytically the risk-neutral process of the underlying returns as a function of the same parameters of the physical process. Thus, the parameters can be estimated by maximum likelihood using underlying returns under the physical measure, and used to simulate returns according to the risk-neutral process in order to price options via Monte Carlo (for a survey and more details, see Christoffersen, Jacobs and Ornthanalai, 2013).

Therefore, the third benchmark we consider is the GARCH option model of Duan (1995). For each day, we estimate the GARCH(1,1) model with normal innovations using daily past underlying returns in a window of four years. To obtain the risk-neutral process, we follow the procedure in Zhu and Ling (2015). Then, for each option with maturity T, we simulate 100,000 paths of T days of daily risk-neutral underlying returns, compound the returns for each path, and price the option via Monte Carlo using the empirical martingale method of Duan and Simonato (1998). In order to compare our method with this approach, we simulate 100,000 paths of T days of daily physical underlying returns, compound the returns for each path, and use them in the estimation of the generalized entropic risk-neutral measures.

The fourth benchmark is the Glosten, Jagannathan and Runkle (1993) (GJR) model, which allows for asymmetry in the GARCH process. This model has been recently used in the option pricing literature (see, for instance, Barone-Adesi, Engle and Mancini, 2008) and has been shown to outperform the GARCH model in Zhu and Ling (2015). We estimate the GJR(1,1) model with normal innovations in a window of four years and follow Zhu and Ling (2015) to obtain the risk-neutral process. Then, for each option with maturity T, we simulate 100,000 paths of T days of daily risk-neutral underlying returns, compound the returns for each path, and price the option via Monte Carlo following Duan and Simonato (1998). In order to compare our method with this approach, we simulate 100,000 paths of T days of daily physical underlying returns, compound the returns for each path, and use them in the estimation of the generalized entropic risk-neutral measures.

The GARCH option pricing methodology relies on the assumption of a particular risk-neutral measure under which it is possible to obtain an analytic risk-neutral process for the underlying. However, in the context of incomplete markets, there is an infinity of admissible measures. Therefore, our method can be seen as providing flexible nonparametric changes of measure for the GARCH process, where option pricing easily follows from equation (2.12). In the next subsection, we investigate if these alternative risk-neutral measures can improve the pricing performance of the GARCH option models.

2.4.2 Empirical Results

We assess the option pricing performance of the generalized entropic estimators and benchmark methods using the MPE and the MAPE as error statistics.[13] We calculate the MPE and MAPE for each moneyness and maturity category for the whole sample of S&P 500 call and put options. For instance, for the category of short-term options with $S/X \in [0.90, 0.97)$, the MPE and MAPE are obtained considering all short-term OTM calls and ITM puts.

For the first empirical exercise, Figure 2.6 plots, for each moneyness and maturity category, the MPE curves over γ obtained by connecting the MPE of each generalized entropic estimator considered. The MPE for the Black-Scholes model is given by a horizontal line. As can be seen, on average, the Black-Scholes model overestimates the prices of OTM calls and ITM puts, and underestimates the prices of OTM puts and ITM calls. This is because the lognormal risk-neutral distribution of the model usually has a fatter right tail than the model-implied risk-neutral distribution, and a thinner left tail.[14] As for the generalized entropic estimators, on average, OTM calls and ITM puts require positive γ's to be priced, especially short-term options. In other words, they require risk-neutral distributions with a thinner right tail. On the other hand, OTM puts and ITM calls require negative γ's between -1 and -0.5 to be priced, on average. Negative γ's are associated to risk-neutral distributions with fatter left tails. Figure 2.7 plots the MAPE for the Black-Scholes model and each generalized entropic estimator considered. For short-term options, the lowest MAPE of OTM calls and ITM puts is achieved with $\gamma = 5$, while for the remaining options the Hellinger estimator performs better. For long-term options, Hellinger and empirical likelihood outperform the remaining estimators, for all moneyness categories.

In the second empirical exercise, we allow for the possibility that there is an ATM option available for the estimation of the implied volatility, as well of the physical volatility. Figure 2.8 reports the MPE results for this exercise. As would be expected, the performance of the Black-Scholes model improves for ATM options, but the biases for far-from-the-money options remain. Again, on average, OTM calls and ITM puts require positive γ's to be priced, while OTM puts and ITM calls require negative γ's. Figure 2.9 plots the MAPE results. Again, positive γ's yield smaller MAPE for low-moneyness options, while γ's close to zero perform better for ATM and high-moneyness options. Note also that the magnitude of the errors is smaller than the one where the historical volatility was used in the estimation.

[13] Results are qualitatively the same under alternative error statistics.

[14] A risk-neutral distribution with a fatter right tail gives higher prices for OTM calls, because there is a higher probability that the options finish ITM. Analogously, a thinner left tail implies lower prices for OTM puts, because there is a lower probability that the options finish ITM.

Overall, the results above show that the Black-Scholes model can be outperformed by most of the generalized entropic estimators considered individually. However, given the incompleteness of the market when using only underlying returns, there is no reason to consider only one estimator. In fact, our analysis identifies the estimators that are appropriate to price different options in the cross-section. Risk-neutral measures coming from positive γ's that decrease the probability mass in large positive returns deliver the smallest pricing errors for OTM calls and ITM puts. As for the remaining options, the best performance is usually given by a γ between -1 and 0, such as the Hellinger estimator. Such estimators are associated to risk-neutral measures that overweight the occurrence of large negative underlying returns.

We now investigate if the nonparametric changes of measure we provide can improve the performance of GARCH options models. The MPE results for the third empirical exercise are reported in Figure 2.10. The GARCH option model presents the same moneyness biases that the Black-Scholes model, but improves the performance for long-term ATM and high-moneyness options. As can be seen, the alternative changes of measure we provide can always improve the performance of the GARCH model. In particular, estimators with positive γ's perform better for low-moneyness options, while negative γ's outperform for high-moneyness options. The same is true under the MAPE results in Figure 2.11 for short-term options. For long-term ATM options and OTM puts and ITM calls, the GARCH option model is only marginally outperformed by the KLIC estimator.

The fourth empirical exercise considers the GJR GARCH option model. From the MPE results in Figure 2.12, we can see that the GJR model improves the performance with respect to both the Black-Scholes model and the GARCH option model. Even so, it still overvalues OTM calls and ITM puts and underestimates the prices of OTM puts and ITM calls. This is such that, again, the generalized entropic estimators can be used to improve the pricing performance. Figure 2.13 shows that the risk-neutral measure coming from $\gamma = 5$ generates considerably smaller MAPE for short-term low-moneyness and ATM options, while the Hellinger estimator outperforms for high-moneyness options. For long-term options, it is harder to improve on the GJR model, but the Euclidean divergence and the KLIC estimators still generate smaller MAPE for low-moneyness and high-moneyness options, respectively.

In summary, we document that benchmark methods estimated using only underlying returns present systematic moneyness biases in option pricing. More specifically, they generate risk-neutral distributions that have thinner left tails and fatter right tails than the option-implied risk-neutral distribution. The same is true if we consider a particular generalized entropic estimator. Nonetheless, we show that we can consider distinct estimators for different options in the cross-section. That is, we can select estimators that correctly compensate for the left or right tails of the risk-neutral distribution separately. We identify that positive γ's are appropriate to price low-moneyness options, while negative γ's such as the Hellinger estimator usually outperform for high-moneyness options.

2.5 Conclusion

In this paper, we address the problem of how to choose appropriate risk-neutral measures using only information from underlying returns in the context of option pricing in incomplete markets. We study the option pricing implications of a family of minimum generalized entropy measures that are the "closest" possible to the physical distribution while being consistent with the underlying returns. Each measure distorts higher moments of the physical probabilities in a particular way and allows for different compensations for the tails of the risk-neutral distribution.

With Monte Carlo experiments, we investigate how the generalized entropic estimators behave under different sources of risk and market incompleteness. In a Black-Scholes environment, we show that there is an optimal estimator that is completely determined by the parameters governing risk premium. That is, when there is only one source of risk coming from the underlying price process, there is a risk-neutral measure in the Cressie-Read family that matches the model-implied risk-neutral distribution. In contrast, in the presence of stochastic volatility and jumps, the optimal estimator for a given option depends not only on the levels of risk premia, but also on the option moneyness and maturity. That is, distinct estimators are necessary to price different options in the cross-section.

In a large-scale empirical application, we compare the out-of-sample option pricing performance of the generalized entropic estimators with benchmark methods. The Black-Scholes model and GARCH option models estimated using only underlying returns present systematic moneyness biases in the cross-section of options. The same is true for our estimators individually. That is, it is not possible to approximate both tails of the option-implied risk-neutral distribution using only information from the physical distribution. Nonetheless, we can select generalized entropic estimators that correctly compensate for the left or right tails of the risk-neutral distribution separately. Therefore, we identify the estimators that are appropriate to price different options in the cross-section.

2.6 Appendix A - Proofs

Proof of Proposition 11

In the Black-Scholes model, we have, under the physical measure:

$$\ln\left(\frac{S_t}{S_u}\right) = \left(\mu - \frac{1}{2}\sigma^2\right)(t-u) + \sigma\left(W_t - W_u\right), \tag{2.14}$$

where μ is the continuous expected rate of return, σ is the volatility and W_t is a Wiener process. Letting $R_{u,t} = \frac{S_t}{S_u}$ and rewriting (2.14), we get:

$$W_t - W_u = \frac{1}{\sigma}\left[\ln\left(R_{u,t}\right) - \left(\mu - \frac{1}{2}\sigma^2\right)(t-u)\right]. \tag{2.15}$$

In order to change to the risk-neutral measure guaranteed to exist by no-arbitrage, one may apply the Girsanov theorem. In this case, the Radon-Nikodym derivative is:

$$Z(t) = \exp\left\{-\theta(W_t - W_0) - \frac{1}{2}\theta^2 t\right\}, \tag{2.16}$$

where $\theta = \frac{\mu - r}{\sigma}$ and r is the risk-free rate. We can write $Z(t)$ as a function of the underlying asset returns:

$$Z(t) = \exp\left\{-\theta\frac{1}{\sigma}\left[\ln\left(R_{0,t}\right) - \left(\mu - \frac{1}{2}\sigma^2\right)t\right] - \frac{1}{2}\theta^2 t\right\},$$

$$\Rightarrow Z(t) = \exp\left\{\ln\left(\left(R_{0,t}\right)^{-\frac{\theta}{\sigma}}\right) + \frac{\theta}{\sigma}\left(\mu - \frac{1}{2}\sigma^2\right)t - \frac{1}{2}\theta^2 t\right\},$$

$$\Rightarrow Z(t) = \left(R_{0,t}\right)^{-\frac{\theta}{\sigma}}\exp\left\{At\right\}, \tag{2.17}$$

where:

$$A = \frac{\theta\mu}{\sigma} - \frac{1}{2}\sigma\theta - \frac{1}{2}\theta^2. \tag{2.18}$$

The fact that the Radon-Nikodym derivative in the Girsanov theorem is a martingale (under the physical measure) with $Z_0 = 1$ implies that $\mathbb{E}\left[Z_t\right] = 1$, directly giving:

$$\mathbb{E}\left[\left(R_{0,t}\right)^{-\frac{\theta}{\sigma}}\right] = \exp\left\{-At\right\}, \tag{2.19}$$

where $\mathbb{E}[\cdot]$ is the expectation under the physical probability measure. Moreover, we have, by the properties of the Radon-Nikodym derivative and risk-neutral measures:

$$\mathbb{E}\left[R_{0,t}Z(t)\right] = \widetilde{\mathbb{E}}\left[R_{0,t}\right] = e^{-rt}, \tag{2.20}$$

where $\widetilde{\mathbb{E}}[.]$ is the expectation under the Black-Scholes risk-neutral measure.

Now, given an option with time to maturity t, define $R = R_{0,t}$. The Cressie-Read minimum discrepancy state-price density for $\gamma < 0$ is given by:

$$\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} = \frac{\left(1 + \gamma\lambda\left(R - R_f\right)\right)^{\frac{1}{\gamma}}}{\mathbb{E}\left[\left(1 + \gamma\lambda\left(R - R_f\right)\right)^{\frac{1}{\gamma}}\right]}. \tag{2.21}$$

Under the Black-Scholes model, R is lognormal (see equation (2.14)). Since the risk-neutral measure has to price the underlying asset, the optimization problem has the following restriction:

$$\frac{1}{R_f}\mathbb{E}^{\mathbb{Q}}\left[R\right] = 1,$$

or, equivalently:

$$\frac{1}{R_f}\mathbb{E}\left[\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}R\right] = 1, \tag{2.22}$$

along with $1 + \gamma\lambda\left(R - R_f\right) > 0$ almost surely. If there is a λ such that (2.22) holds and $1+\gamma\lambda\left(R - R_f\right) > 0$, then $\mathrm{d}\mathbb{Q}/\mathrm{d}\mathbb{P}$ from (2.21) will be the solution. Moreover, this solution will be unique because the dual problem is strictly concave. Now, define (implicitly) $\widehat{\lambda}$ as $\lambda = \frac{\widehat{\lambda}}{\gamma R_f}$ to obtain:

$$\begin{aligned}
\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} &= \frac{\left(1 + \gamma\frac{\widehat{\lambda}}{\gamma R_f}\left(R - R_f\right)\right)^{\frac{1}{\gamma}}}{\mathbb{E}\left[\left(1 + \gamma\frac{\widehat{\lambda}}{\gamma R_f}\left(R - R_f\right)\right)^{\frac{1}{\gamma}}\right]}, \\
\Rightarrow \frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} &= \frac{\left(\frac{1}{R_f}\right)^{\frac{1}{\gamma}}\left(R_f + \widehat{\lambda}\left(R - R_f\right)\right)^{\frac{1}{\gamma}}}{\left(\frac{1}{R_f}\right)^{\frac{1}{\gamma}}\mathbb{E}\left[\left(R_f + \widehat{\lambda}\left(R - R_f\right)\right)^{\frac{1}{\gamma}}\right]}, \\
\Rightarrow \frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} &= \frac{\left(R_f(1-\widehat{\lambda}) + \widehat{\lambda}R\right)^{\frac{1}{\gamma}}}{\mathbb{E}\left[\left(R_f(1-\widehat{\lambda}) + \widehat{\lambda}R\right)^{\frac{1}{\gamma}}\right]}.
\end{aligned} \tag{2.23}$$

By making the ansatz $\widehat{\lambda} = 1$, R_f is eliminated from the risk-neutral measure:

$$\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} = \frac{R^{\frac{1}{\gamma}}}{\mathbb{E}\left[R^{\frac{1}{\gamma}}\right]}.$$

Now, in order to find the appropriate γ, just compare the above $\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}$ with $Z(t)$. The comparison suggests the choice $1/\gamma^* = -\theta/\sigma$, which makes $\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} = Z(t)$, where t is the time to maturity and $\mathbb{E}\left[R^{\frac{1}{\gamma}}\right] = \exp\left\{-At\right\}$ by (2.19). In order to show that this is the solution, it remains to verify that $1 + \gamma\lambda\left(R - R_f\right) > 0$ and that (2.22) holds. Indeed,

$1+\gamma\lambda\left(R-R_f\right)=R/R_f>0$ almost surely and, because R is lognormal and noting that $R_f=e^{-rt}$, we have:

$$\frac{1}{R_f}\mathbb{E}\left[\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}R\right]=\frac{1}{e^{-rt}}\mathbb{E}\left[Z(t)R\right]=1,$$

as expected.

2.7 Tables

Table 2.1: **Parameters for the SVCJ model**

Measure	μ or r	κ_v	θ_v	σ_v	ρ	λ	μ_s	σ_s	μ_v	ρ_s
$\mathbb{P}$	0.1396	6.5520	0.0135	0.08	-0.4838	1.5120	-0.0263	0.0289	0.0373	0.0
$\mathbb{Q}_1$	0.05	14.112	0.0062	0.08	-0.4838	1.5120	-0.0658	0.0289	0.2724	0.0
$\mathbb{Q}_2$	0.05	14.364	0.0061	0.08	-0.4838	1.5120	-0.0539	0.0578	0.2213	0.0
$\mathbb{Q}_3$	0.05	6.5520	0.0135	0.08	-0.4838	1.5120	-0.0725	0.0289	0.1333	0.0
$\mathbb{Q}_4$	0.05	6.5520	0.0135	0.08	-0.4838	1.5120	-0.0501	0.0751	0.0935	0.0

This table presents the annualized parameters used in the simulations of the SVCJ model. The parameters for the physical measure ($\mathbb{P}$) and risk-neutral measure ($\mathbb{Q}$) correspond to the ones estimated in Eraker, Johannes and Polson (2003) and Broadie, Chernov and Johannes (2007), respectively. Four specifications for the risk-neutral measure are considered. $\mathbb{Q}_1$ and $\mathbb{Q}_2$ allow for $\kappa_v^{\mathbb{Q}}\neq\kappa_v$, while $\mathbb{Q}_3$ and $\mathbb{Q}_4$ impose $\kappa_v^{\mathbb{Q}}=\kappa_v$. $\mathbb{Q}_1$ and $\mathbb{Q}_3$ impose $\sigma_s^{\mathbb{Q}}=\sigma_s$, while $\mathbb{Q}_2$ and $\mathbb{Q}_4$ allow for $\sigma_s^{\mathbb{Q}}\neq\sigma_s$.

Table 2.2: **Summary Statistics of S&P 500 Index Options**

	Call Options			Put Options		
Moneyness	Short	Long	Subtotal	Short	Long	Subtotal
	\$8.19	\$28.91	\$15.31	\$106.60	\$133.19	\$114.78
$[0.90, 0.97)$	13.11%	14.26%	13.51%	16.35%	16.73%	16.47%
	{162148}	{84959}	{247107}	{53070}	{23597}	{76667}
	\$34.96	\$75.10	\$42.10	\$38.46	\$77.22	\$45.73
$[0.97, 1.03)$	13.17%	16.31%	13.73%	14.02%	16.73%	14.53%
	{373929}	{80864}	{454793}	{364781}	{84261}	{449042}
	\$113.01	\$148.67	\$120.89	\$13.49	\$45.79	\$19.73
$[1.03, 1.10]$	18.83%	19.26%	18.93%	18.43%	19.08%	18.56%
	{90970}	{25796}	{116766}	{381304}	{91416}	{472720}
	\$39.36	\$64.52	\$45.25	\$31.07	\$69.43	\$38.73
Subtotal	13.98%	15.80%	14.40%	16.28%	17.81%	16.58%
	{627047}	{191619}	{818666}	{799155}	{199274}	{998429}

This table presents summary statistics of the S&P 500 index option data after applying our filtering to the OptionMetrics Database. The sample ranges from January 4, 1996 to June 28, 2019. The columns Short and Long refer to the maturity categories. For each moneyness (S/X) and maturity category, the first row depicts the average option price, the second row the average implied volatility and the third row the number of observations (in braces). The average of the daily values of the S&P 500 index and the (annualized) risk-free rate in the sample period were 1443.61 and 2.17%, respectively.

2.8 Figures

Figure 2.1: **MPE - Black-Scholes Environment**

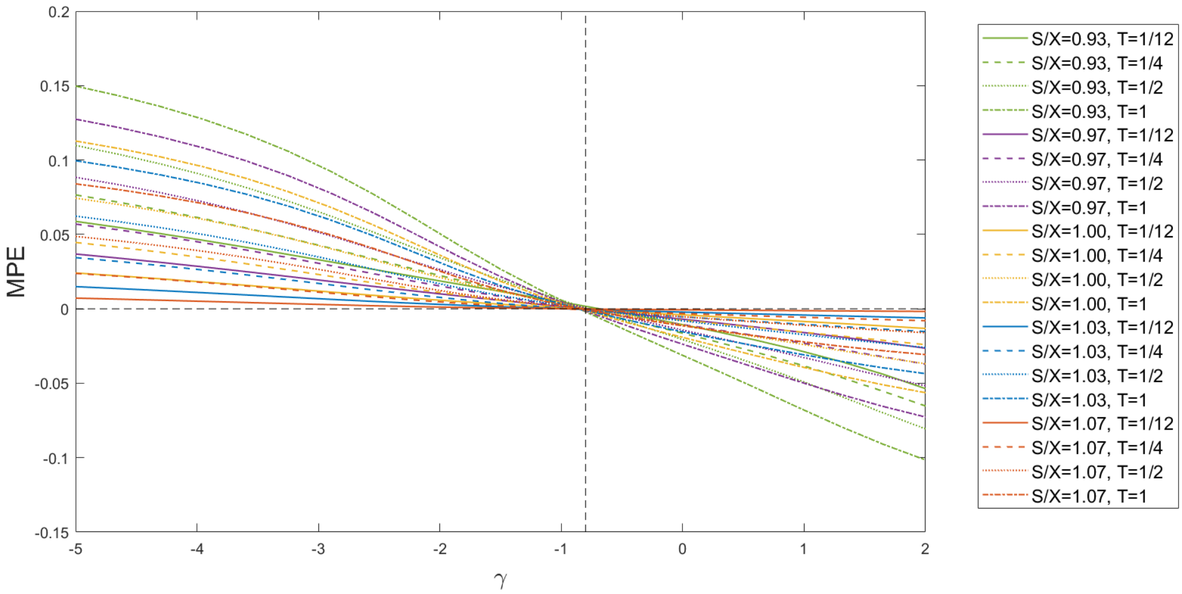

This figure plots the MPE for a range of γ's in the Black-Scholes environment for European call options with different combinations of moneyness and maturity. The vertical dashed line corresponds to $\gamma = -0.8$. For each γ, the MPE of a given option is calculated as the average across 5000 simulations of the percentage pricing error of the generalized entropic estimator using 200 returns from the physical distribution of the model. S/X represents the option moneyness, and T is the time to maturity in years.

Figure 2.2: **MPE Surface - Black-Scholes Environment**

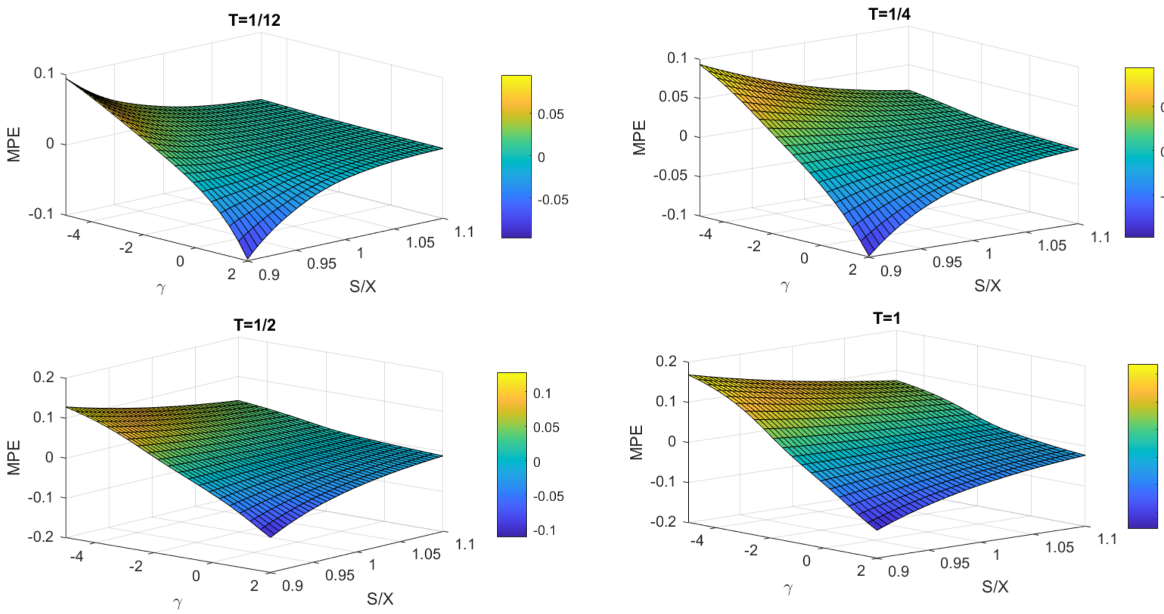

This figure plots the MPE surface for a range of γ's and moneynesses (S/X) in the Black-Scholes environment for different maturities (T, in years). For each γ, the MPE of a given option is calculated as the average across 5000 simulations of the percentage pricing error of the generalized entropic estimator using 200 returns from the physical distribution of the model.

Figure 2.3: **MAPE Surface - Black-Scholes Environment**

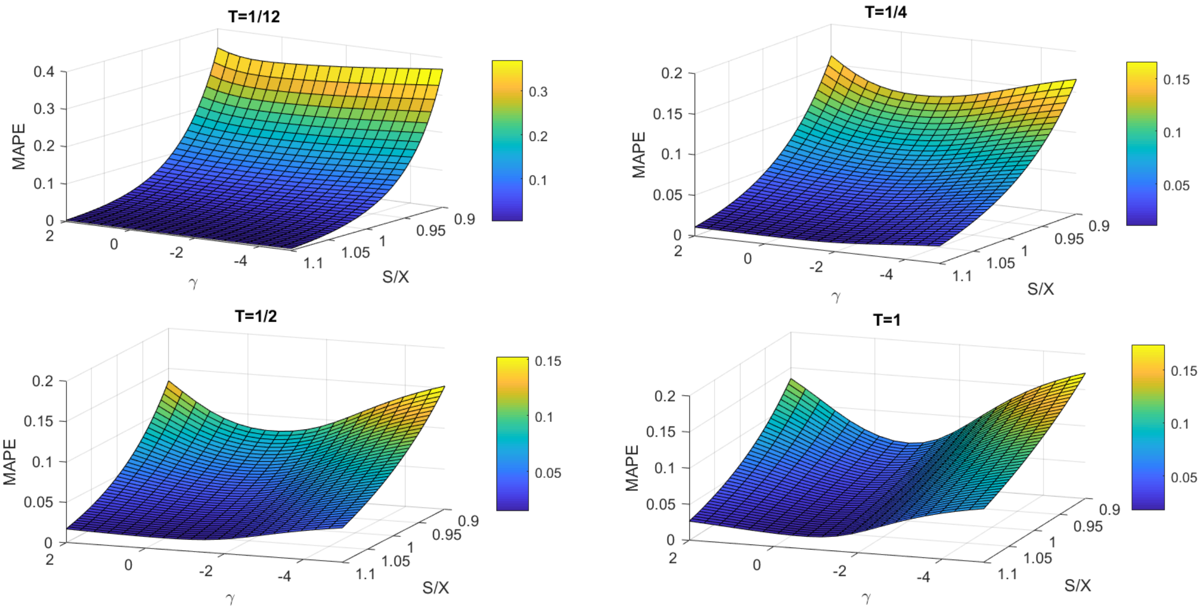

This figure plots the MAPE surface for a range of γ's and moneynesses (S/X) in the Black-Scholes environment for different maturities (T, in years). For each γ, the MAPE of a given option is calculated as the average across 5000 simulations of the absolute percentage pricing error of the generalized entropic estimator using 200 returns from the physical distribution of the model.

Figure 2.4: **MPE - SVCJ Environment**

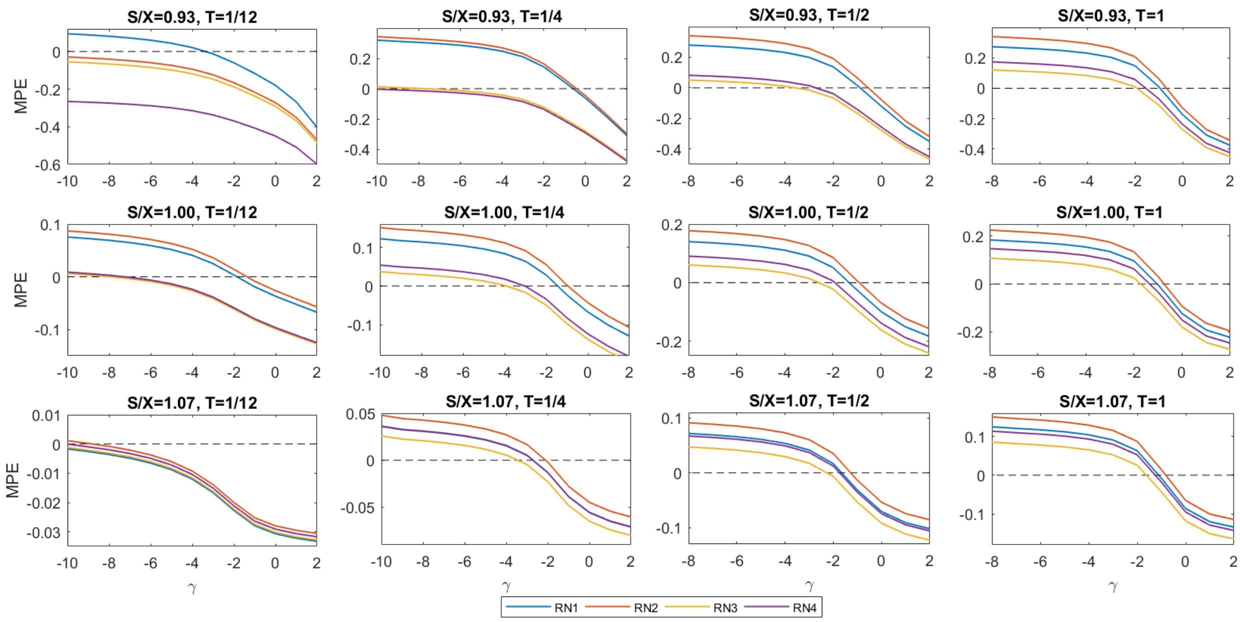

This figure plots the MPE for a range of γ's and four different risk-neutral distributions in the SVCJ environment for European call options with different combinations of moneyness and maturity. Each risk-neutral distribution (RN1-RN4) gives a different theoretical option price as benchmark. For each γ and risk-neutral distribution, the MPE of a given option is calculated as the average across 5000 simulations of the percentage pricing error of the generalized entropic estimator using 200 returns from the physical distribution of the model. S/X represents the option moneyness, and T is the time to maturity in years.

Figure 2.5: **MAPE Surface - SVCJ Environment**

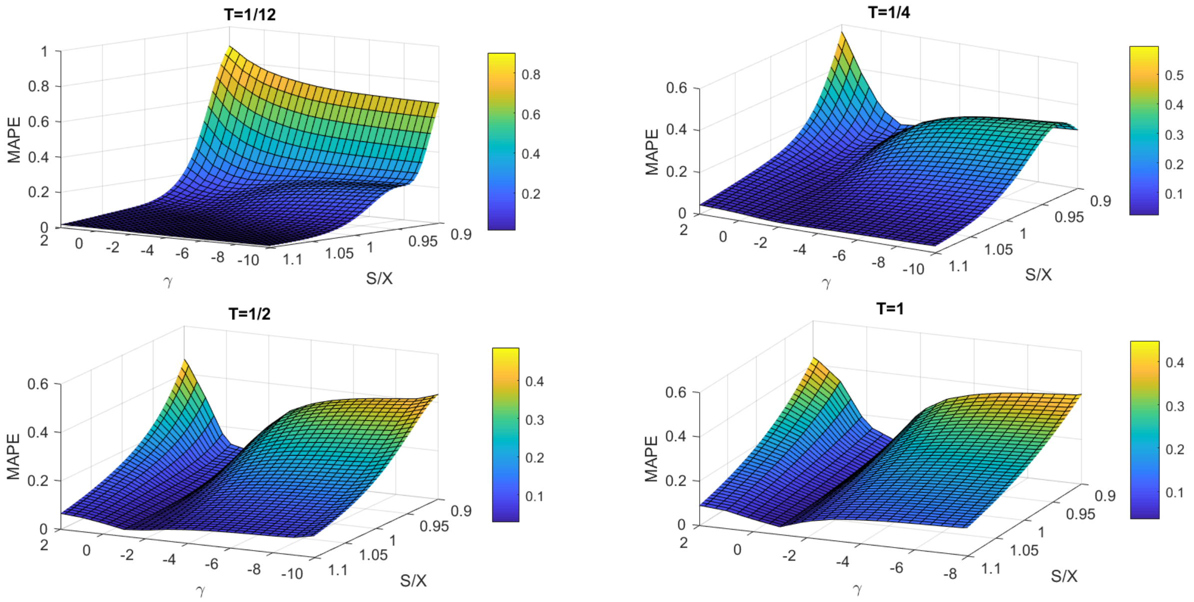

This figure plots the MAPE surface for a range of γ's and moneynesses (S/X) in the SVCJ environment for different maturities (T, in years), considering the risk-neutral parametrization $\mathbb{Q}_2$. For each γ, the MAPE of a given option is calculated as the average across 5000 simulations of the percentage absolute pricing error of the generalized entropic estimator using 200 returns from the physical distribution of the model.

Figure 2.6: **MPE - Generalized Entropic Estimators and Black-Scholes with Historical Volatility**

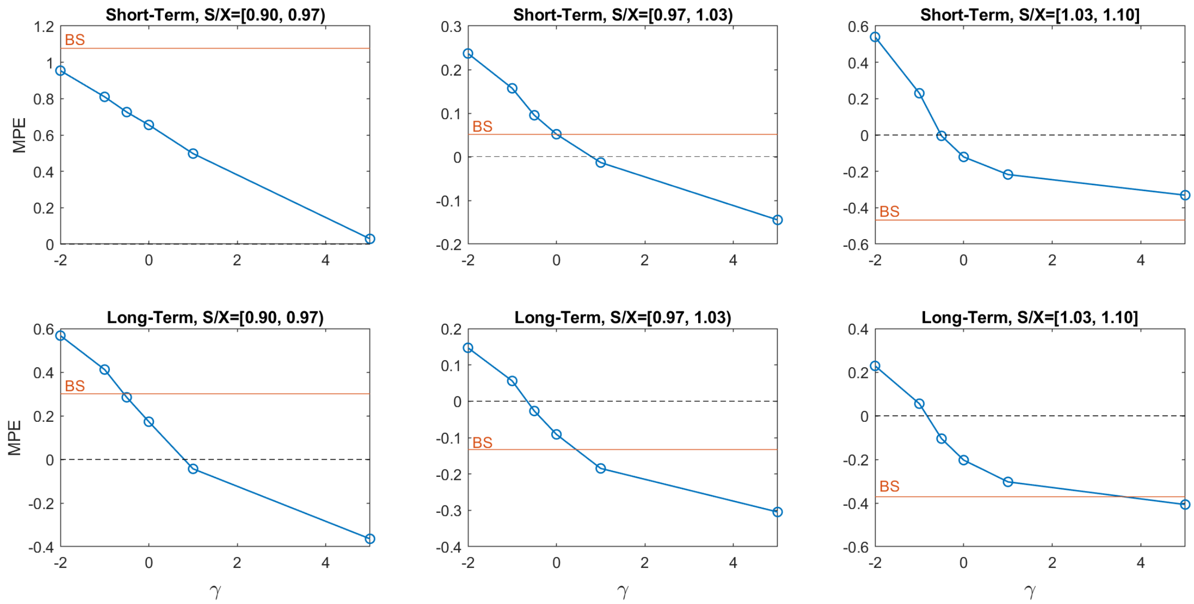

This figure plots the MPE for each moneyness and maturity category of the generalized entropic estimators and the Black-Scholes (BS) model estimated with historical volatility. The sample ranges from January 4, 1996 to June 28, 2019.

Figure 2.7: **MAPE - Generalized Entropic Estimators and Black-Scholes with Historical Volatility**

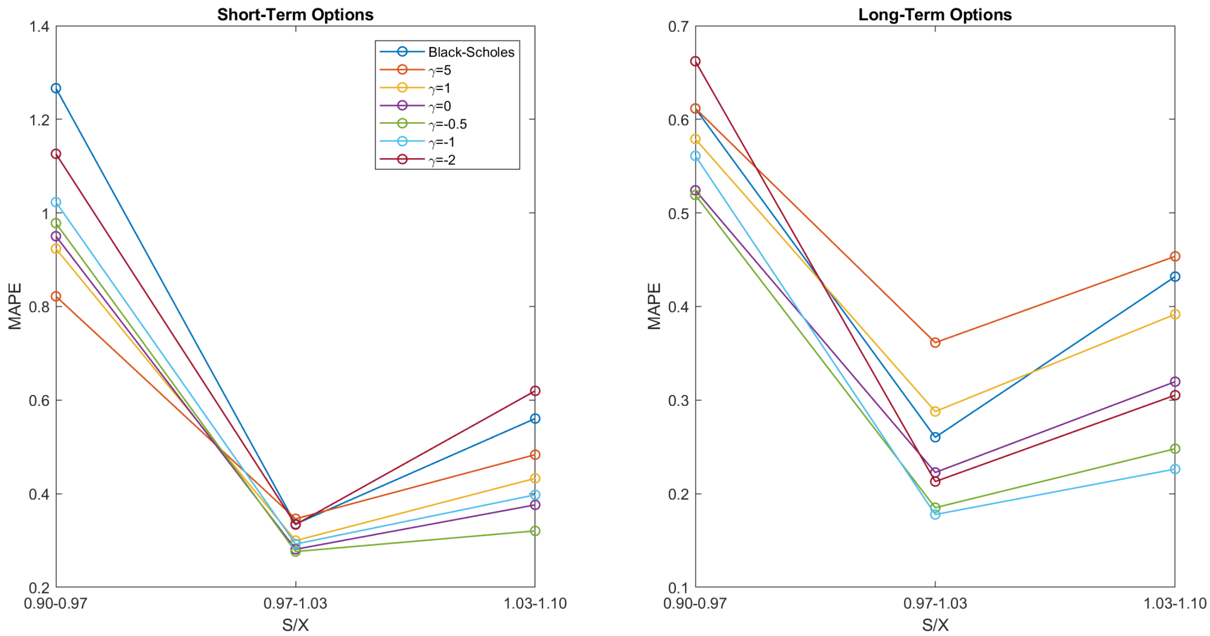

This figure plots the MAPE for each moneyness and maturity category of the generalized entropic estimators and the Black-Scholes model estimated with historical volatility. The sample ranges from January 4, 1996 to June 28, 2019.

Figure 2.8: **MPE - Generalized Entropic Estimators and Black-Scholes with ATM Implied Volatility**

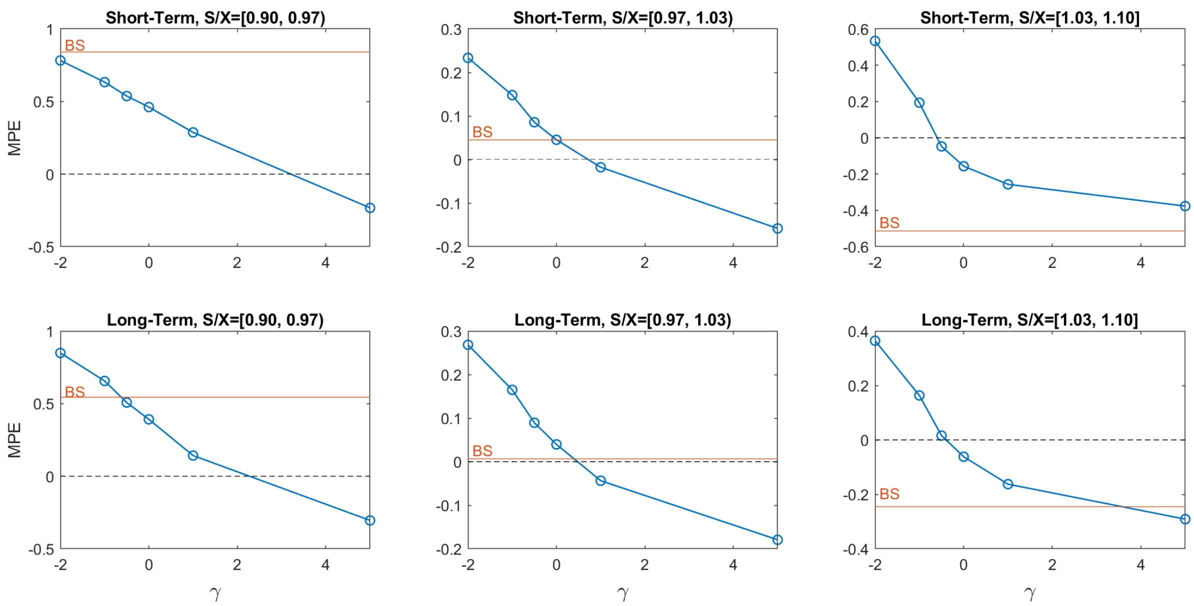

This figure plots the MPE for each moneyness and maturity category of the generalized entropic estimators and the Black-Scholes (BS) model estimated with ATM implied volatility. The sample ranges from January 4, 1996 to June 28, 2019.

Figure 2.9: **MAPE - Generalized Entropic Estimators and Black-Scholes with ATM Implied Volatility**

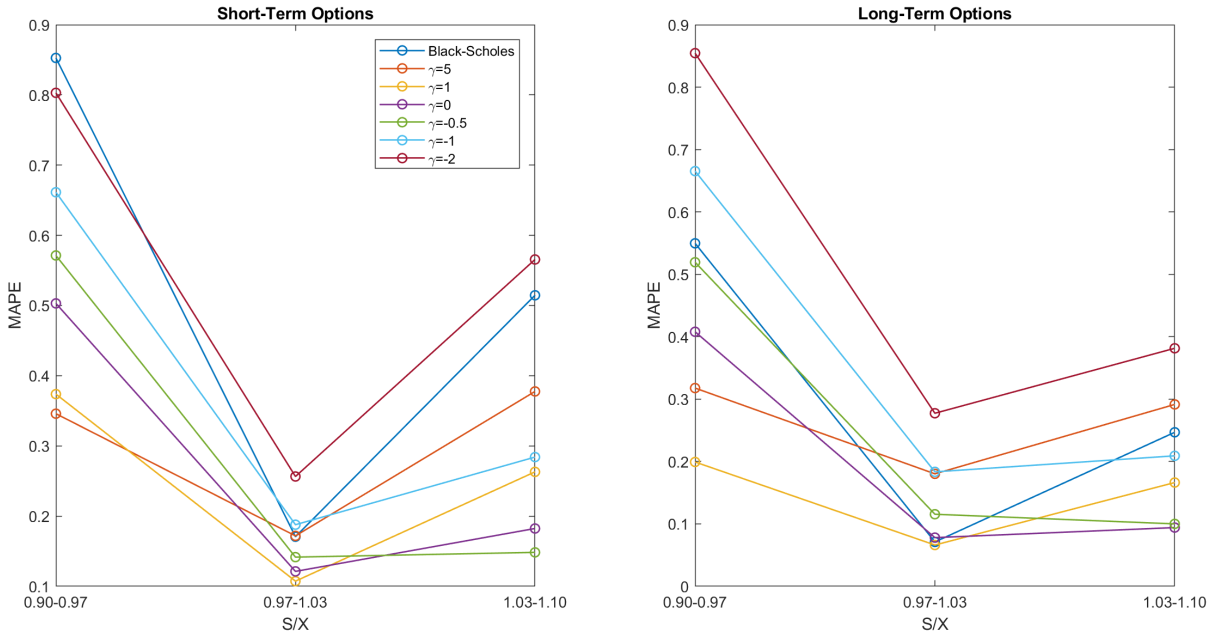

This figure plots the MAPE for each moneyness and maturity category of the generalized entropic estimators and the Black-Scholes model estimated with ATM implied volatility. The sample ranges from January 4, 1996 to June 28, 2019.

Figure 2.10: **MPE - Generalized Entropic Estimators and GARCH**

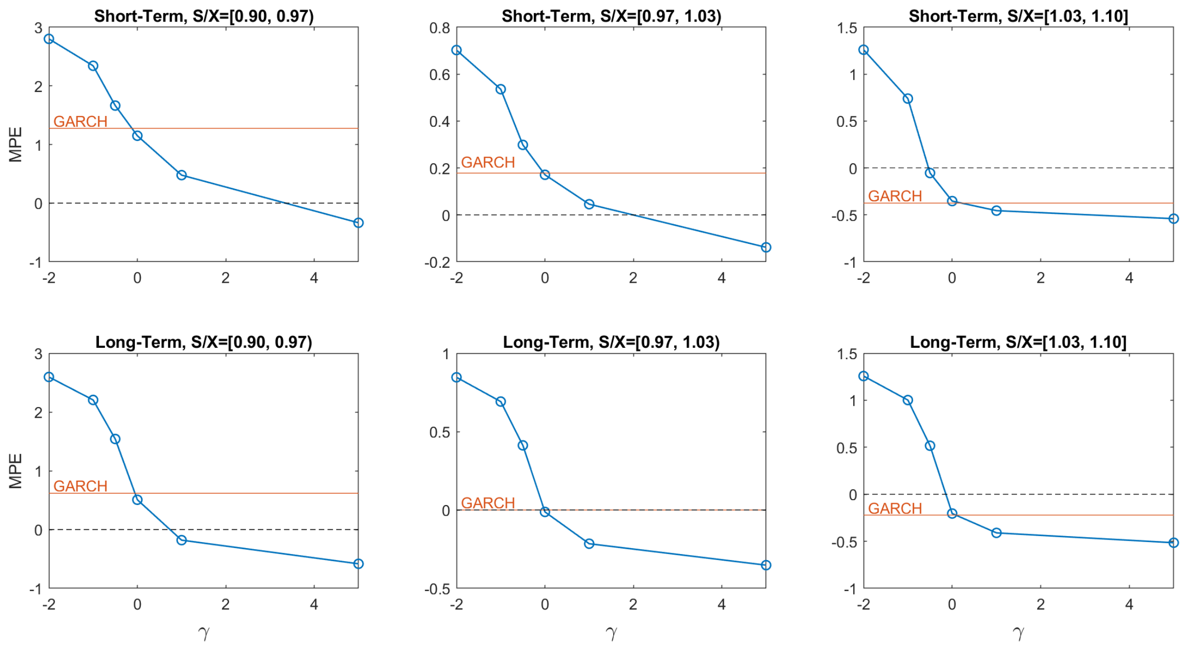

This figure plots the MPE for each moneyness and maturity category of the generalized entropic estimators and the GARCH option model. The sample ranges from January 4, 1996 to June 28, 2019.

Figure 2.11: **MAPE - Generalized Entropic Estimators and GARCH**

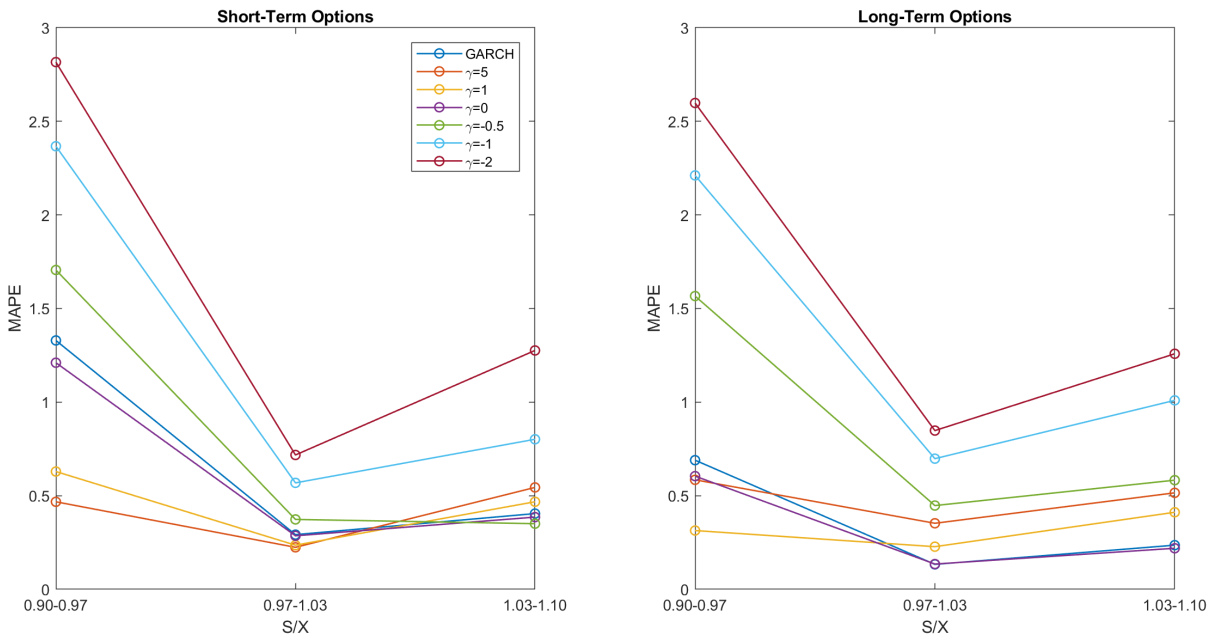

This figure plots the MAPE for each moneyness and maturity category of the generalized entropic estimators and the GARCH option model. The sample ranges from January 4, 1996 to June 28, 2019.

Figure 2.12: **MPE - Generalized Entropic Estimators and GJR GARCH**

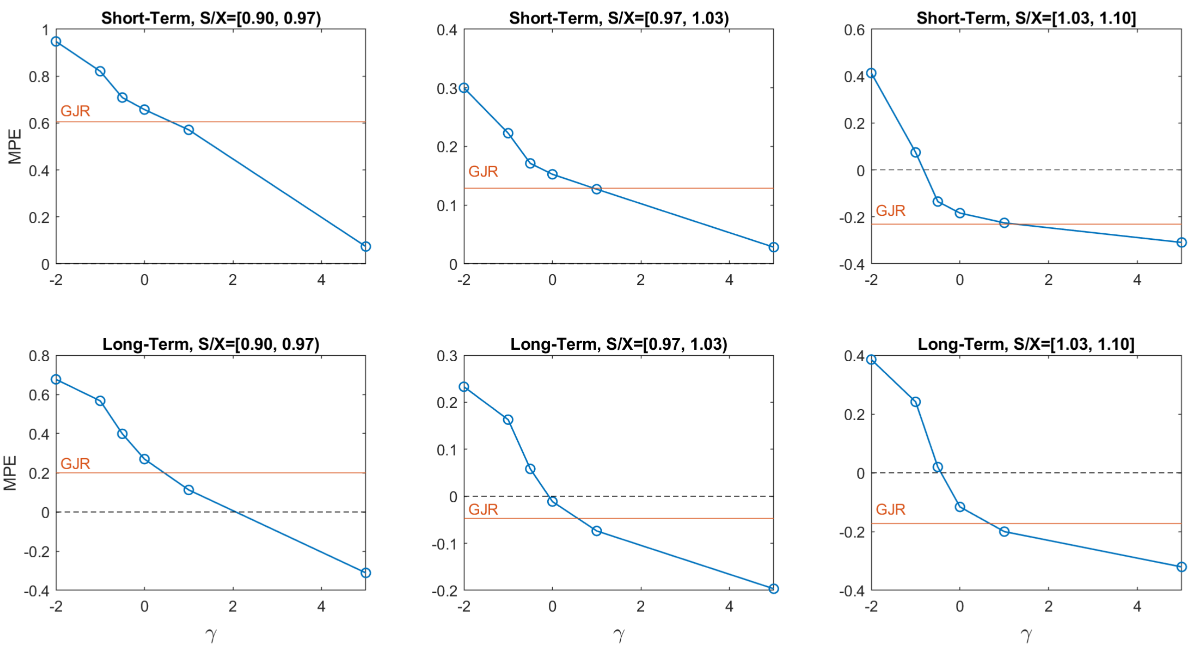

This figure plots the MPE for each moneyness and maturity category of the generalized entropic estimators and the GJR GARCH option model. The sample ranges from January 4, 1996 to June 28, 2019.

Figure 2.13: **MAPE - Generalized Entropic Estimators and GJR GARCH**

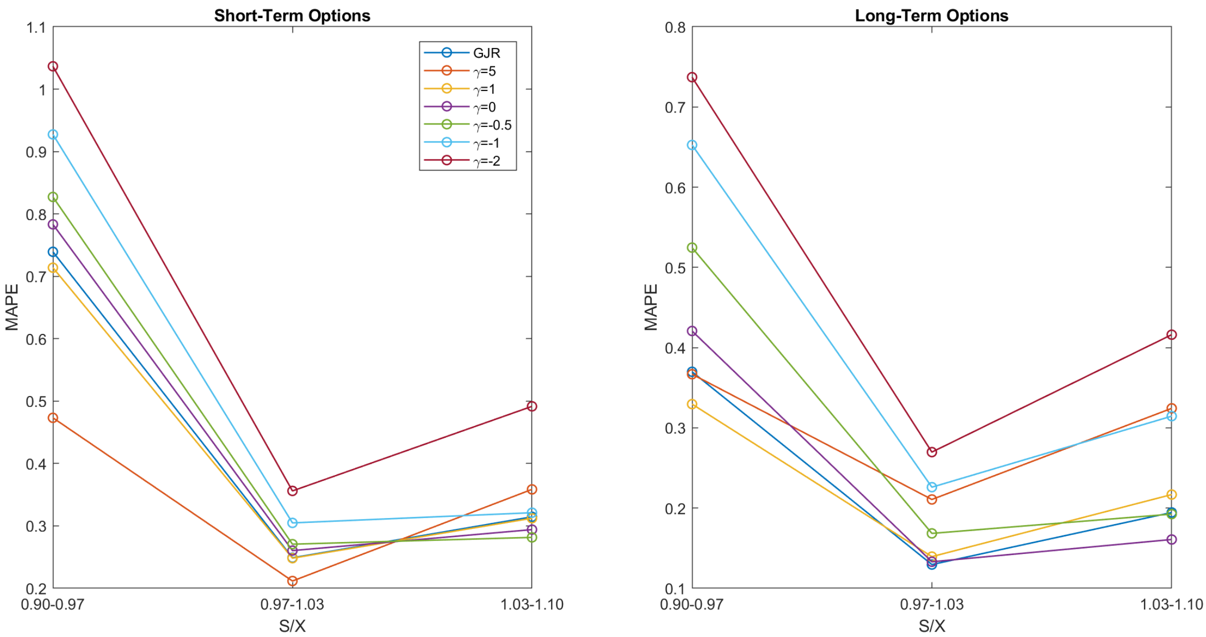

This figure plots the MAPE for each moneyness and maturity category of the generalized entropic estimators and the GJR GARCH option model. The sample ranges from January 4, 1996 to June 28, 2019.

Chapter 3

Endogenous Wealth and the Pricing Kernel Puzzle

We provide a unifying framework for the literature that calculates empirical pricing kernels (EPKs) as the ratio of the option-implied state-price density and the historical return distribution of the market index. Such EPKs are often a U-shaped function of market returns, which is puzzling under the assumption of a complete market where the index proxies for wealth. We propose to estimate the pricing kernel by minimizing a convex discrepancy function subject to correctly pricing the underlying asset and observed options. The projection of our estimated pricing kernel onto market returns is the EPK identified by the literature. However, by duality, we are also able to obtain the pricing kernel as a function of the endogenous wealth of the marginal investor pricing the index and index options. This implies that the U-shaped EPK can be rationalized as the projection of a pricing kernel that is monotonically decreasing in the endogenous wealth. That is, there is no puzzle when we recognize that options also constitute investment opportunities. The U-shaped pattern of the EPK arises because the marginal investor in the index and index options sells protection against large movements in the index.

This chapter is co-authored with Caio Almeida.

3.1 Introduction

In the absence of arbitrage, the current price of any asset is given by the expectation over different states of nature of the future asset payoff multiplied by the pricing kernel. Thus, the pricing kernel, also known as the stochastic discount factor (SDF), embeds all information about investors' risk preferences that is relevant for pricing. Given its utmost importance, there is a large literature devoted to estimate and study the SDF. In particular, beginning with Ait-Sahalia and Lo (2000) and Jackwerth (2000), several papers have used index option data and historical market returns to nonparametrically estimate an "empirical projection" of the economy-wide SDF onto the space of index returns.[1] The empirical pricing kernel (EPK) is calculated as the ratio of the state-price density (SPD), estimated from the cross-section of index option prices, and the physical distribution, estimated from the time series of market returns.[2]

Ait-Sahalia and Lo (2000) and Jackwerth (2000) assume a complete market framework with a representative investor where the market index is perfectly correlated with aggregate wealth. Under these assumptions, the projection of the pricing kernel onto market returns is equal to the economy-wide SDF. Empirically, they obtain EPKs that are not a monotonically decreasing function of wealth (as proxied by the S&P 500 index). In fact, EPKs are usually U-shaped in times of high market volatility, and W-shaped or tilde-shaped in times of low volatility (Cuesdeanu, 2016). The violation of monotonicity is inconsistent with a risk-averse investor, which has been labeled as the pricing kernel puzzle by Brown and Jackwerth (2012).

In this paper, we argue that the assumptions mentioned above, which are followed by most of the literature on the pricing kernel puzzle, are inconsistent with the estimation procedure of the EPK. To show that, we provide a unifying framework that directly connects the estimation of the SPD to the estimation of the pricing kernel. We estimate the SPD that minimizes a convex loss function measuring the dispersion with respect to the physical distribution, subject to correctly pricing observed options and the underlying asset. This procedure is very similar to the usual nonparametric method to recover risk-neutral probabilities (Rubinstein, 1994; Jackwerth and Rubinstein, 1996; Jackwerth, 2000).[3] In fact, it leads to the same SPD given a sufficiently large cross-section of options. The advantage, however, is that our method is equivalent to a dual problem akin to the maximization of a concave utility function, as in Almeida and Garcia (2017). In particular,

---

[1] See, for instance, Jackwerth (2004), Beare and Schmidt (2014), Cuesdeanu (2016) and Song and Xiu (2016).

[2] Other papers use option data and historical returns to parametrically estimate the EPK. See Rosenberg and Engle (2002), Barone-Adesi, Engle and Mancini (2008) and Christoffersen, Heston and Jacobs (2013).

[3] The difference is that these papers consider lognormal prior probabilities instead of the empirically observed probabilities characterizing the physical distribution.

the pricing kernel associated to the SPD comes from an investor choosing the portfolio that maximizes expected wealth among the observed options and market index. Our method permits to recover the optimal portfolio weights and the resulting endogenous wealth. This allows us to obtain the pricing kernel not only as a function of the market index, but also of the endogenous wealth taking into account all the investment opportunities of the investor.

We investigate the implications of our framework to the pricing kernel puzzle in simulated economies coming from well-known parametric models. The first economy is a Black and Scholes (1973) environment where there is no puzzle, as the EPK is monotonically decreasing with respect to the market index. Using our method, we first estimate the SPD and pricing kernel using only market returns. The estimated SPD exactly matches the model-implied SPD, while the pricing kernel projection onto market returns is equal to the EPK. This suggests that options are redundant in this economy. We confirm this by including options in the estimation of the SPD and pricing kernel. The associated investor sets zero weights to the options in the portfolio, only trading in the market index. Moreover, we show that the endogenous wealth coming from the portfolio optimization equals the market index. Therefore, there is no pricing kernel puzzle because options are redundant and the assumption that the market index proxies for aggregate wealth is valid.

The second economy incorporates additional sources of risk given by stochastic volatility and jumps (Bates, 2000; Duffie, Pan and Singleton, 2000). In this economy, the pricing kernel puzzle exists as the risk-neutral distribution is more skewed to the left and has fatter tails than the physical distribution, generating a U-shaped EPK. In contrast to the Black-Scholes economy, options are non-redundant and are necessary in the estimation in order to approximate the model-implied SPD and EPK. We show that the U-shaped EPK is the projection onto market returns of the estimated pricing kernel that is monotonically decreasing over the endogenous wealth that considers the market index and index options as investment opportunities. That is, there is no puzzle as the investor is risk-averse. The U-shaped pattern arises because extreme negative and positive market returns are associated to negative returns on the optimal portfolio, leading to smaller wealth. This illustrates that the assumption of perfect correlation between the market index and wealth is misspecified in the presence of stochastic volatility and jumps risk premia. Moreover, it indicates that the marginal investor in the index and index options sells protection against large movements in the index.

The remaining of the paper is organized as follows. Section 3.2 presents the pricing kernel puzzle. Section 3.3 describes our framework for the estimation of the SPD and pricing kernel. Section 3.4 investigates the implications of our method to the pricing kernel puzzle in simulated economies. Section 3.5 discusses how our results relate to the literature on the pricing kernel puzzle and concludes the paper.

3.2 The Pricing Kernel Puzzle

In the absence of arbitrage, the current price P_t of any asset is given by the expectation of the future asset payoff X_T multiplied by the pricing kernel $m_{t,T}$:[4]

$$P_t = \mathbb{E}^{\mathbb{P}}[m_{t,T} X_T | \mathcal{F}_t] \equiv \int_0^\infty X_T(s)\, m_{t,T}(s)\, \pi_{t,T}^{\mathbb{P}}(s)\, ds, \tag{3.1}$$

where s represents the state of the economy, $\mathcal{F}_t$ is the information available to investors at time t and $\pi_{t,T}^{\mathbb{P}}$ is the probability density function (PDF) under the physical measure $\mathbb{P}$. Also known as the stochastic discount factor (SDF), the pricing kernel distorts the physical measure in such a way that one can take simple expectations to calculate the price of any asset. More specifically, given the almost sure positivity of $m_{t,T}$ under no-arbitrage, the pricing kernel induces a change of measure from the physical measure $\mathbb{P}$ to the risk-neutral measure $\mathbb{Q}$. Given a risk-free rate R_f from t to T, this can be seen by noting that $\mathbb{E}_t[m_{t,T}] = 1/R_f$, and dividing and multiplying (3.1) by $\mathbb{E}_t[m_{t,T}]$:

$$P_t = \frac{1}{R_f} \int_0^\infty X_T(s) \frac{m_{t,T}(s)}{\mathbb{E}_t[m_{t,T}]} \pi_{t,T}^{\mathbb{P}}(s)\, ds = \frac{1}{R_f} \int_0^\infty X_T(s)\, \pi_{t,T}^{\mathbb{Q}}(s)\, ds \equiv \frac{1}{R_f} \mathbb{E}^{\mathbb{Q}}[X_T | \mathcal{F}_t], \tag{3.2}$$

where $\pi_{t,T}^{\mathbb{Q}}$ is the PDF under the risk-neutral measure $\mathbb{Q}$. This PDF is known as the state-price density (SPD), as it defines the prices of Arrow-Debreu securities paying one dollar at time T if state of nature s is realized, and zero elsewhere. From (3.1) and (3.2), it becomes evident that the pricing kernel is the ratio of discounted risk-neutral probabilities ($\pi_{t,T}^{\mathbb{Q}}/R_f$) and physical probabilities ($\pi_{t,T}^{\mathbb{P}}$):

$$m_{t,T} = \frac{\pi_{t,T}^{\mathbb{Q}}}{R_f \pi_{t,T}^{\mathbb{P}}}. \tag{3.3}$$

Ait-Sahalia and Lo (2000) and Jackwerth (2000) empirically derive the SDF based on the equation above. They separately estimate $\pi_{t,T}^{\mathbb{Q}}$ and $\pi_{t,T}^{\mathbb{P}}$, and then calculate the empirical pricing kernel (EPK) as the ratio between them. The risk-neutral probabilities are estimated from the cross-section of S&P 500 option prices, while the physical probabilities are estimated from the time series of S&P 500 returns.

The first papers showing how to nonparametrically estimate the SPD from observed option prices were Rubinstein (1994) and Jackwerth and Rubinstein (1996). They choose the risk-neutral probabilities that minimize a loss function with respect to prior

[4] See Cochrane (2001) for a comprehensive treatment of the pricing kernel.

lognormal probabilities and are consistent with the underlying asset prices and the prices of observed options. Jackwerth (2000) and Jackwerth (2004) adopt variations of this method designed to be faster and more stable, but the intuition remains the same. Other papers, such as Ait-Sahalia and Lo (1998) and Ait-Sahalia and Lo (2000), obtain the SPD by using the Breeden and Litzenberger (1978) formula, which relates the second derivative of option prices with respect to the strike to the risk-neutral distribution.

For the estimation of the physical probabilities, Ait-Sahalia and Lo (2000) and Jackwerth (2000) use a kernel density to smooth the histogram of realized market returns over previous years. Kernel densities only rely on the assumption of stationarity of the returns. On the other hand, Rosenberg and Engle (2002) and Barone-Adesi, Engle and Mancini (2008) obtain conditional estimates for the physical distribution taking into account time-varying volatility. They fit the Glosten, Jagannathan and Runkle (1993) GARCH model to historical returns.

It is important to note that the EPK estimated as described above is not necessarily equal to the economy-wide SDF. In traditional consumption-based asset pricing models (e.g. Lucas, 1978), the SDF equals the intertemporal marginal rate of substitution between consumption in t and T. More generally, however, the pricing kernel can be a function of multiple state variables. The EPK avoids the issue of specifying the state variables, being defined over realizations of the market payoff. As Rosenberg and Engle (2002) note, the EPK can be interpreted as the projection of the economy-wide SDF onto market returns. This projected pricing kernel has the same pricing implications than the original SDF for assets that depend on the market payoff.

Ait-Sahalia and Lo (2000) and Jackwerth (2000) assume a complete market with a representative investor where the index level is perfectly correlated with aggregate wealth. Under these assumptions, the projected pricing kernel onto market returns is equal to the economy-wide SDF, which in turn is proportional to the marginal utility of the representative investor. Empirically, these papers find that the EPK exhibits non-decreasing parts with respect to wealth (as proxied by the S&P 500 index level). This implies that the representative investor is locally risk-seeking, contradicting standard asset pricing theory. The violation of monotonicity with respect to wealth characterizes the pricing kernel puzzle (Brown and Jackwerth, 2012).

3.3 A Unifying Framework for the Estimation of the Pricing Kernel

In this section, we provide a unifying framework that, given the physical distribution, directly connects the estimation of the SPD to the estimation of the pricing kernel.

Under this framework, it becomes clear that the nonparametric estimation of the SPD and the pricing kernel (and, as a consequence, of the EPK) is equivalent, by duality, to the maximization of a concave utility function. In particular, the EPK is the projection onto market returns of the estimated pricing kernel that is proportional to the marginal utility of an investor trading in the market index and the index options. Our method allows to recover the optimal portfolio weights and the resulting endogenous wealth of the investor for each state of nature. This is important as we are able to obtain the pricing kernel not only as a function of the market index, but also of the endogenous wealth that considers all the investment opportunities of the investor.

3.3.1 Theoretical Background

Suppose we observe at time t the market index S_t and the prices of $K-1$ index options expiring at time T. An admissible SDF will be a random variable $m_{t,T}$ satisfying:

$$\mathbb{E}^{\mathbb{P}}[m_{t,T}S_T|\mathcal{F}_t] = S_t, \tag{3.4}$$

$$\mathbb{E}^{\mathbb{P}}[m_{t,T}g_j(S_T)|\mathcal{F}_t] = C_{j,t}, \quad j = 1, ..., K-1, \tag{3.5}$$

where S_T is a random variable representing the market index at time T, $g_j(S_T)$ is the payoff function of option j and $C_{j,t}$ is the price of option j.[5] It will be useful to work with the pricing equations above in terms of excess returns. We stack the return on the market S_T/S_t and the returns on the options $g_j(S_T)/C_{j,t}$ on a K-dimensional random vector $\mathbf{R}$.[6] Denote by $\mathbf{R}^e \equiv \mathbf{R} - \mathbf{R}_f$ the vector of excess returns with respect to the risk-free rate, where $\mathbf{R}_f \equiv R_f \mathbf{1}_K$ and $\mathbf{1}_K$ is a conformable vector of ones. Then, (3.4) and (3.5) can be succinctly written as:

$$\mathbb{E}[m\mathbf{R}^e] \equiv \int m\mathbf{R}^e \,\mathrm{d}\mathbb{P} = \mathbf{0}_K. \tag{3.6}$$

In the absence of arbitrage, there exists a strictly positive admissible SDF, which in turn defines a risk-neutral measure via the change of measure $\mathrm{d}\mathbb{Q} = \frac{m}{\mathbb{E}[m]}\,\mathrm{d}\mathbb{P}$:

$$\int \mathbf{R}^e \frac{m}{\mathbb{E}[m]}\,\mathrm{d}\mathbb{P} = \int \mathbf{R}^e \,\mathrm{d}\mathbb{Q} \equiv \mathbb{E}^{\mathbb{Q}}[\mathbf{R}^e] = \mathbf{0}_K. \tag{3.7}$$

Under the restriction that $\mathbb{Q}$ is absolutely continuous with respect to $\mathbb{P}$, $\mathrm{d}\mathbb{Q}/\mathrm{d}\mathbb{P}$ is a Radon-Nikodym derivative.[7]

---

[5] For a call option with strike Y_j, $g_j(S_T) = max(S_T - Y_j, 0)$, while for a put option with strike Y_j, $g_j(S_T) = max(Y_j - S_T, 0)$.

[6] As the analysis goes through in a one-period model, we drop time subscripts from now on.

[7] That is, $\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} : \Omega \to [0,\infty)$ is a measurable function such that for any measurable set $A \subseteq \Omega$, $\mathbb{Q}(A) = \int_A \frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\mathrm{d}\mathbb{P}$.

We propose to choose the risk-neutral probabilities that minimize a convex loss function with respect to the physical distribution and are consistent with the market index and index option prices. Letting $\phi(.)$ denote a general convex function, we solve the following minimization in the space of admissible measures with $I^{\phi}(\mathbb{Q},\mathbb{P}) \equiv \mathbb{E}\left[\phi\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)\right] < \infty$:

$$\mathbb{Q}^* = arg\ \underset{Q}{min}\ \mathbb{E}\left[\phi\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)\right] \equiv \int \phi\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)\mathrm{d}\mathbb{P}, \quad \text{s.t. } \mathbb{E}^{\mathbb{Q}}\left[\mathbf{R}^e\right] = \mathbf{0}_K, \tag{3.8}$$

where the risk-neutral measure must also be nonnegative and integrate to one. The problem above is equivalent to estimating the SDF that minimizes dispersion according to the convex loss function $\phi(.)$, as in Almeida and Garcia (2017). Letting m^* denote the minimum dispersion SDF, the following relationship holds:

$$m^* = \frac{1}{R_f}\frac{\mathrm{d}\mathbb{Q}^*}{\mathrm{d}\mathbb{P}}. \tag{3.9}$$

While problem (3.8) is of infinite dimension and difficult to solve, in Chapter 1 we show, using the results in Almeida and Garcia (2017), that it can be solved via a much simpler finite dimensional dual problem. The only required condition for the primal and dual problems to coincide is the absence of arbitrage. The dual problem consists of the maximization of a convex conjugate that is akin to a concave utility function:

$$\underset{\lambda\in R^K}{max}\ -\mathbb{E}\left[\phi^{*,+}(1+\lambda'\mathbf{R}^e) + \delta(\lambda \,|\, \Lambda(\mathbf{R}^e))\right], \tag{3.10}$$

where $\Lambda(\mathbf{R}^e) = \{\lambda \in \mathbb{R}^K : (1+\lambda'\mathbf{R}^e) \in \text{dom } \phi^{*,+}\}$,[8] $\delta(.\,|\,C)$ is such that $\delta(x\,|\,C) = 0$ if $x \in C$ and ∞ otherwise, and $\phi^{*,+}$ denotes the convex conjugate of ϕ.[9] The vector λ is composed of the Lagrange multipliers coming from the pricing equations in (3.8) for the market index and the index options.

The economic interpretation of the dual problem is better appreciated when we specify the convex loss function $\phi(.)$. We follow Almeida and Garcia (2017) and Chapter 1 in considering the Cressie and Read (1984) family of convex loss functions:

$$\phi_\gamma\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right) = \frac{\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)^{\gamma+1} - 1}{\gamma(\gamma+1)}, \quad \gamma \in \mathbb{R}. \tag{3.11}$$

In Chapter 1 we show that, for each loss function in the Cressie-Read family as indexed by γ, the dual problem coming from the corresponding minimization in (3.8) can be economically interpreted as a standard optimal portfolio problem for an investor with

[8] The domain of $\phi^{*,+}(z)$ is defined as the values of z for which the function is finite ($\phi^{*,+}(z) < \infty$).
[9] $\phi^{*,+}(z) = \underset{w\in[0,\infty)\cap\text{domain }\phi}{sup} zw - \phi(w)$.

HARA utility:

$$u^{\gamma}(W) = -\frac{1}{\gamma+1}(b - a\gamma W)^{\frac{\gamma+1}{\gamma}}, \tag{3.12}$$

with $a > 0$ and $b - a\gamma W > 0$, which guarantees that the function u^{γ} is well-defined, concave and strictly increasing. More specifically, the investor distributes her initial wealth W_0 putting $\tilde{\lambda}_j$ units of wealth on the risky asset R_j and the remaining $W_0 - \sum_{j=1}^{K} \tilde{\lambda}_j$ in a risk-free asset paying R_f. Terminal wealth is then given by $W(\tilde{\lambda}) = W_0 R_f + \sum_{j=1}^{K} \tilde{\lambda}_j (R_j - R_f)$ and the investor maximizes expected utility:

$$\tilde{\lambda}_{\gamma}^{*} = \underset{\tilde{\lambda} \in R^K}{max}\ \mathbb{E}\left[u^{\gamma}(W(\tilde{\lambda}))\right]. \tag{3.13}$$

Solving (3.13) is equivalent to solving the dual problem (3.10) when the convex loss function in the primal problem is a member of the Cressie-Read family. In particular, letting λ_{γ}^{*} denote the dual problem solution, $\tilde{\lambda}_{\gamma}^{*} = -\lambda_{\gamma}^{*}(b - a\gamma W_0 R_f)/a$ if $\gamma \neq 0$ and $\tilde{\lambda}_{\gamma}^{*} = -\lambda_{\gamma}^{*}/a$ if $\gamma = 0$. This result essentially shows that the minimum dispersion SDFs are proportional to the marginal utilities of HARA investors. The HARA class encompasses as special cases several widely used utility functions. If $\gamma = 1$, it specializes to quadratic utility, while $\gamma \to 0$ yields exponential utility. The power utility function is obtained if $\gamma < 0$ and $a = -1/\gamma$, specializing to CRRA utility if, further, $b = 0$. Logarithmic utility obtains for $a = 1$ and $\gamma \to -1$, while linear utility (the risk-neutral case) is achieved when $\gamma \to -\infty$.

The framework above sheds light on the estimation and interpretation of the EPK. Nonparametrically estimating the SPD consistent with the market index and index option prices is equivalent to estimating the pricing kernel correctly pricing these assets. By duality, the estimated pricing kernel is proportional to the marginal utility of an investor trading in the index and the options in order to maximize expected utility over wealth. Given that the utility function is concave, this SDF will be monotonically decreasing in the endogenous optimal wealth resulting from the portfolio optimization. The EPK is the projection of the estimated SDF onto the returns on the market index. The advantage of our method is that we are able to obtain not only this projection, but also the SDF as a function of the endogenous wealth, which is consistent with all the investment opportunities of the investor.

3.3.2 Estimation Procedure

In order to describe the estimation procedure in our method, it is convenient to consider a finite state space framework, with $k = 1, ..., n$ states of nature, where $n \geq K$. Let $\{\mathbf{R}_k\}_{k=1}^{n}$ be the observed gross returns of the K risky assets (the market index and the index options), where each $\mathbf{R}_k$ is independent and identically distributed according to $\mathbb{P}$. The unknown physical measure $\mathbb{P}$ can be replaced by the empirical measure $\mathbb{P}_n$ that gives weights $\pi_k = 1/n$ to each draw from the physical distribution, which can be obtained from the historical market returns. This allows us to exchange the expectation $\mathbb{E}^{\mathbb{P}}$ with its sample counterpart $\sum_{k=1}^{n} \pi_k^{\mathbb{P}} \equiv \frac{1}{n} \sum_{k=1}^{n}$. The state-price density will be given by the discrete risk-neutral probabilities $\pi_k^{\mathbb{Q}}$. In the following corollary, we summarize the results from Chapter 1 for the sample version of the problem of finding a minimum dispersion Cressie-Read risk-neutral measure:

Corollary 6. *Consider the primal problem:*

$$\begin{aligned} &\underset{\{\pi_1^Q,\ldots,\pi_n^Q\}}{min} \quad \textstyle\sum_{k=1}^{n} \pi_k^{\mathbb{P}} \frac{\left(\pi_k^{\mathbb{Q}}/\pi_k^{\mathbb{P}}\right)^{\gamma+1}-1}{\gamma(\gamma+1)}, \\ &s.t. \;\; \textstyle\sum_{k=1}^{n} \pi_k^{\mathbb{Q}} \left(\boldsymbol{R}_k - \boldsymbol{R}_f\right) = \boldsymbol{0}_K, \;\; \sum_{k=1}^{n} \pi_k^{\mathbb{Q}} = 1, \; \pi_k^{\mathbb{Q}} \geq 0 \;\; \forall\, k. \end{aligned} \tag{3.14}$$

Absence of arbitrage in the observed sample implies that the value of the primal problem coincides (with dual attainment) with the value of the dual problem below:

i) if $\gamma > 0$*:*

$$\lambda_\gamma^* = \underset{\lambda \in R^K}{arg\ max} - \frac{1}{\gamma+1} \sum_{k=1}^{n} \pi_k^{\mathbb{P}} \left(1 + \gamma\lambda' \left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)\right)^{\left(\frac{\gamma+1}{\gamma}\right)} I_{\Lambda_\gamma(\boldsymbol{R}_k)}(\lambda), \tag{3.15}$$

ii) if $\gamma < 0$*:*

$$\lambda_\gamma^* = \underset{\lambda \in \Lambda_\gamma}{arg\ max} - \frac{1}{\gamma+1} \sum_{k=1}^{n} \pi_k^{\mathbb{P}} \left(1 + \gamma\lambda' \left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)\right)^{\left(\frac{\gamma+1}{\gamma}\right)}, \tag{3.16}$$

iii) if $\gamma = 0$*, the maximization is unconstrained:*

$$\lambda_0^* = \underset{\lambda \in R^K}{arg\ max} \; - \sum_{k=1}^{n} \pi_k^{\mathbb{P}} e^{\lambda'\left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)}, \tag{3.17}$$

where $\Lambda_\gamma = \{\lambda \in \mathbb{R}^K : for\ k = 1, ..., n,\ (1 + \gamma\lambda' (\boldsymbol{R}_k - \boldsymbol{R}_f)) > 0\}$*,* $\Lambda_\gamma(\boldsymbol{R}_k) = \{\lambda \in \mathbb{R}^K : (1 + \gamma\lambda' (\boldsymbol{R}_k - \boldsymbol{R}_f)) > 0\}$ *and* $I_A(x) = 1$ *if* $x \in A$*, and* 0 *otherwise.*

The minimum dispersion risk-neutral measure can then be recovered from the first-

order conditions of (3.15), (3.16) and (3.17) with respect to λ, evaluated at λ_γ^*:

$$\pi_k^{\mathbb{Q}*}(\gamma, \boldsymbol{R}) = \frac{(1 + \gamma\lambda_\gamma^{*\prime}(\boldsymbol{R}_k - \boldsymbol{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\boldsymbol{R}_k)}(\lambda_\gamma^*)}{\sum_{i=1}^n (1 + \gamma\lambda_\gamma^{*\prime}(\boldsymbol{R}_i - \boldsymbol{R}_f))^{\frac{1}{\gamma}} I_{\Lambda_\gamma(\boldsymbol{R}_i)}(\lambda_\gamma^*)}, \quad k = 1, ..., n; \ \gamma > 0, \tag{3.18}$$

$$\pi_k^{\mathbb{Q}*}(\gamma, \boldsymbol{R}) = \frac{\left(1 + \gamma\lambda_\gamma^{*\prime}\left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)\right)^{\frac{1}{\gamma}}}{\sum_{i=1}^n \left(1 + \gamma\lambda_\gamma^{*\prime}\left(\boldsymbol{R}_i - \boldsymbol{R}_f\right)\right)^{\frac{1}{\gamma}}}, \ k = 1, ..., n; \ \gamma < 0, \tag{3.19}$$

$$\pi_k^{\mathbb{Q}*}(0, \boldsymbol{R}) = \frac{e^{\lambda_0^*\left(\boldsymbol{R}_k - \boldsymbol{R}_f\right)}}{\sum_{i=1}^n e^{\lambda_0^*\left(\boldsymbol{R}_i - \boldsymbol{R}_f\right)}}, \ k = 1, ..., n; \ \gamma = 0. \tag{3.20}$$

Proof. See Chapter 1. □

The results above show how to obtain, for a given γ, the risk-neutral probability (and the pricing kernel, by dividing by $R_f \pi_k^{\mathbb{P}}$) consistent with the market index and the index option prices for each state of nature. The minimum dispersion risk-neutral measure is a nonlinear function of the excess returns on the optimal portfolio of risky assets, being defined over the realizations of these returns. In particular, from the estimated Lagrange multipliers λ_γ^*, for a given initial wealth W_0 we can recover the portfolio weights of the investor as $\tilde{\lambda}_\gamma^* = -\lambda_\gamma^*(b - a\gamma W_0 R_f)/a$ and the endogenous wealth as $W_k = W_0 R_f + \sum_{j=1}^K \tilde{\lambda}_j (R_{k,j} - R_f)$. Therefore, we can project the SDF not only as a function of the returns of each risky asset, but also with respect to the endogenous wealth. It is also possible to assess how different the endogenous wealth is from the market index, which is often considered as a proxy for wealth.

When applying the method, one has to choose the loss function measuring dispersion in the Cressie-Read family, i.e., the parameter γ. Each loss function generates a risk-neutral measure that is a different hyperbolic function of the optimal portfolio returns, being consistent with the marginal utility of a particular HARA investor. For $\gamma \leq 0$, the minimum dispersion measures are strictly positive, while for $\gamma > 0$ there can be zeros in some states of nature due to the indicator function in (3.18). Moreover, the measures are convex for $\gamma < 1$, concave for $\gamma > 1$ and linear for the case that $\gamma = 1$ (for more details, see Chapter 1).

In the context of using option prices in the estimation, it turns out that the more complete the set of options across strikes, the less sensitive the risk-neutral probabilities will be to the loss function.[10] At the extreme with a continuum of options across strikes, the constraints completely determine the solution, as options payoffs are known to approximate the true SPD (see Breeden and Litzenberger, 1978). While minimum dispersion risk-neutral measures with different γ's can have different shapes, their histograms across market returns states (which approximate the SPD) will get more similar to the

[10] Jackwerth and Rubinstein (1996) also point this out in their similar method.

extent that more options are considered in the estimation. Therefore, given a sufficiently large cross-section of options, different loss functions will produce approximately the same risk-neutral distributions.

3.4 Simulation Evidence

We investigate the implications of our framework to the pricing kernel puzzle with simulated economies coming from well-known parametric models. With the simulations, we can isolate and understand the effects of different sources of risk and market incompleteness on the puzzle. The parametric models allow for analytic solutions to the option pricing equation that can be used to compute the model-implied risk-neutral distribution via the Breeden and Litzenberger (1978) formula. They also allow us to sample underlying returns directly from the model-implied physical distribution.

The first economy considered is a Black and Scholes (1973) environment, while the second economy comes from the stochastic volatility and correlated jumps (SVCJ) model (Bates, 2000; Duffie, Pan and Singleton, 2000). The SVCJ model captures important empirical stylized facts in equity markets and reliable estimates of its physical and risk-neutral parameters have been obtained in the option pricing literature. See Appendix B 1.8 of Chapter 1 for details on the models, the parametrization adopted and the simulation procedure.

Without loss of generality, we set the initial value for the market index to $S_t = 1$. This makes it equivalent to refer to the future market index S_T or to the gross return on the market S_T/S_t. We consider a time horizon of one month, which means that we use one-month underlying returns and prices of options with one month to expiration. For the estimation of the minimum dispersion risk-neutral measure, we use 200,000 simulated one-month underlying returns from the model-implied physical distribution.

3.4.1 Black-Scholes Economy

In a Black-Scholes economy, Rubinstein (1976) shows that the Black-Scholes formula for the price of an option can be equivalently obtained via the pricing kernel of a representative investor with CRRA utility function trading in the market index.[11] This implies that the EPK in this economy should be a nonlinear monotonically decreasing

[11]The Black-Scholes formula is traditionally derived under the assumption of dynamic trading, where the underlying stock and bond can be continuously traded to perfectly hedge an option payoff. That is, dynamic trading completes the market. Rubinstein (1976) shows that it is still possible to obtain the Black-Scholes formula with discrete trading by considering CRRA preferences, i.e., the preferences complete the market.

function of market returns. This is confirmed in Figure 3.1, where we calculate the risk-neutral and physical distributions as the lognormal densities with the corresponding risk-neutral and physical parameters. The distributions are plotted in the left panel, while the ratio between them, characterizing the EPK, is depicted in the right panel. As can be seen, there is no pricing kernel puzzle in this economy, as the EPK is monotonically decreasing with respect to the market index, considered to be a proxy for aggregate wealth.

Our framework allows to obtain the same EPK resulting from the Black-Scholes model using only market returns in the estimation. More specifically, since CRRA is a particular case of the HARA class of utility functions, there is a minimum dispersion risk-neutral measure (and corresponding pricing kernel) that is consistent with the CRRA representative investor. This measure is completely determined by the parameters governing the physical distribution of returns (for more details, see Chapter 1). In the parametrization considered, the risk-neutral measure is given by $\gamma^* = -0.8$. Using the simulated underlying returns from the physical distribution, we calculate the risk-neutral measure and obtain the corresponding SPD via Breeden and Litzenberger (1978). The left panel of Figure 3.2 confirms that the estimated risk-neutral distribution matches the model-implied SPD, while the right panel shows that the estimated pricing kernel projected onto market returns is equal to the EPK.[12] Therefore, in a Black-Scholes economy, the true SPD and EPK can be obtained without using any option data. In other words, options are redundant with respect to the underlying index.

The fact that options are redundant with respect to the index begs the question of what happens if we include options in the estimation of the SPD and the pricing kernel. To investigate that, we include the prices of five options in the estimation of the minimum dispersion risk-neutral measure with $\gamma^* = -0.8$.[13] We find that the investor sets zero weights to the options in the portfolio. In fact, the investor buys approximately the same amount of the market index as the investor trading only on the market. This confirms the redundancy of options in this economy. In the left panel of Figure 3.3, we plot the SDFs estimated using only market returns and using market returns and options, over the endogenous wealth.[14] As can be seen, they are basically the same. Moreover, the right panel shows that the SDF estimated with options is also equal to the model-implied

[12]Note that the estimated pricing kernel is discrete and defined for each realization of state of nature as given by each simulated market return from the physical distribution.

[13]More specifically, we consider three out-of-the-money (OTM) puts with strike prices 0.8, 0.89 and 0.99 and two OTM calls with strikes 1.09 and 1.20. The results are robust to other specifications. Using only OTM options is interesting as OTM puts only pay off for negative realizations of the market, while OTM calls only pay off for positive realizations. This allows for a better interpretation of the optimal portfolio of the investor, while for the purposes of estimating the SPD it is equivalent to using only calls or puts for all strikes.

[14]We set $W_0 = 1$ in order to make it comparable with the market index, and $b = 0$ and $a = -1/\gamma$ in order to obtain the endogenous wealth of the CRRA investor.

EPK when projected onto market returns. More than that, by comparing the two plots, we can see that the endogenous wealth is equal to the market index.

In sum, there is no pricing kernel puzzle in the Black-Scholes economy. Using our framework, we show that this is because options are redundant and the endogenous wealth of the representative investor equals the market index, as assumed by the literature on the puzzle.

3.4.2 Stochastic Volatility and Correlated Jumps Economy

The SVCJ model presents additional sources of risk and market incompleteness given by stochastic volatility and jumps, generating a more realistic economy capable of reproducing stylized facts of real option markets.[15] Using option prices implied by the model, we calculate the risk-neutral distribution using the Breeden and Litzenberger (1978) formula. For the physical probabilities, we derive the kernel density with a Gaussian kernel from one million underlying returns sampled from the model-implied physical distribution.[16] As depicted in the left panel of Figure 3.4, the risk-neutral distribution is more skewed to the left, has fatter tails and is more peaked than the physical distribution. This generates the U-shaped EPK in the right panel of the figure. Under the assumptions of Ait-Sahalia and Lo (2000) and Jackwerth (2000), the violation of monotonicity with respect to wealth (as proxied by the market index) indicates that the representative investor is locally risk-seeking in the regions where the EPK is increasing. Therefore, the pricing kernel puzzle exists in the SVCJ economy.

We start by investigating the implications of using only market returns on the estimation of the SPD and the pricing kernel. Figure 3.5 reports the results for the minimum dispersion risk-neutral measure (and pricing kernel) with $\gamma = -1$. The estimated SPD is clearly not a good approximation to the model-implied SPD, indicating that it is necessary to include options in the estimation of the SPD. Consequently, the projection of the minimum dispersion SDF onto market returns is also very different from the EPK, as can be seen in the right panel of the figure. In particular, it is monotonically decreasing with respect to the market index. This is because the index is perfectly correlated with the endogenous wealth of the investor trading only on the market.

We proceed by including a set of eleven option prices in the estimation of the minimum dispersion SPD and pricing kernel with $\gamma = -1$.[17] The results are reported in Figure 3.6. From the left panel, we can see that the estimated SPD closely matches the

[15]The market generated by the SVCJ model is incomplete even with dynamic trading.

[16]We follow Jackwerth (2000) in setting the kernel bandwidth to $h = 1.8\sigma/n^{1/5}$, where σ is the standard deviation of the sample returns and n is the number of observations.

[17]More specifically, we consider five OTM puts with strike prices 0.8, 0.84, 0.88, 0.92, 0.96 and six OTM calls with strikes 1.01, 1.04, 1.08, 1.12, 1.16 and 1.19. The results are robust to other specifications.

model-implied SPD. Likewise, the projection onto market returns of the estimated SDF approximates the model-implied EPK. This confirms that options are non-redundant with respect to the index in the SVCJ economy.

Our framework highlights that the estimation of the pricing kernel using option prices is inherently related to the investor whose investment opportunities include the market index and the index options. In other words, the market index is only a good proxy for wealth if options are redundant, as in the Black-Scholes economy. Otherwise, the shape of the EPK should not be evidence in favor or against economic theory. This can be seen in Figure 3.7, where we plot the SDF estimated with market returns and option prices over endogenous wealth in the left panel, and its projection onto market returns in the right panel.[18] The U-shaped projection is rationalized by a monotonically decreasing pricing kernel with respect to the endogenous wealth coming from the optimal portfolio of the investor. That is, the investor is risk-averse.

To help visualize the connection between the two plots in Figure 3.7, each color represents the same set of realizations of states of nature. As can be seen, the EPK is U-shaped because extreme negative and positive realizations of the market are associated to negative realizations of the optimal portfolio, leading to smaller wealth. In contrast, moderate positive market returns are the ones associated to larger endogenous wealth. This illustrates how the assumption of perfect correlation between the market index and wealth can be misspecified. More than that, it indicates that the marginal investor in the index and index options sells protection against large movements in the index. This can also be seen by analyzing the investor portfolio weights. In our estimation, the investor sells the OTM put with strike 0.92 and the OTM call with strike 1.12. Due to these positions, the investor incurs in substantial losses for extreme negative and positive returns on the market.

In summary, we show that the U-shaped EPK is not puzzling because it is the projection of a pricing kernel that is monotonically decreasing with respect to the endogenous wealth that considers all available investment opportunities. In particular, the assumption that the index is perfectly correlated with aggregate wealth is misspecified and inconsistent with the estimation procedure of the EPK. The U-shaped pattern of the EPK arises due to the fact that options are non-redundant in the presence of volatility and jumps risk premia, where the marginal investor in the index and index options sells protection against large movements in the index.

[18]We set $W_0 = 1$ in order to make it comparable with the market index, and $b = 0$ and $a = 1$ in order to obtain the endogenous wealth of an investor with logarithmic utility. However, the results are qualitatively the same for any other parametrization.

3.5 Discussion and Final Comments

There is a large literature trying to explain the pricing kernel puzzle, which is surveyed by Cuesdeanu and Jackwerth (2018). We provide a brief summary of the literature and discuss how our findings relate to it. Ziegler (2007) considers a complete market with multiple investors, finding that neither aggregation of heterogeneous preferences, misestimation of beliefs, nor heterogeneous beliefs can reasonably explain the puzzle. Hens and Reichlin (2013) examine violations of the assumptions of risk-averse behavior, unbiased beliefs and complete markets with simple stylized examples. They find that risk-seeking behavior and biased beliefs cannot plausibly account for the puzzle. On the other hand, market incompleteness in the form of background risk can generate the puzzle.

One way of specifying background risk is to introduce additional state variables. Chabi-Yo, Garcia and Renault (2008) show that the puzzle can be explained by regime-switches in a latent state variable. Conditional on the state variable, the pricing kernel is monotonically decreasing in the index, while its projection onto market returns can be U-shaped. Brown and Jackwerth (2012) introduce volatility as an additional state variable. For each volatility state, the pricing kernel is monotonically decreasing, while the SDF given by the expectation over the states can violate monotonicity. Christoffersen, Heston and Jacobs (2013) develop a GARCH option model with a pricing kernel allowing for a variance risk premium. While the pricing kernel is monotonic in the market return and in variance, its projection onto market returns is U-shaped when the variance premium is negative.

While the papers above assume that the market index proxies for aggregate wealth, we show that the nonparametric estimation of the pricing kernel using the market index and index options is inherently related to the marginal investor trading in both markets. Thus, the pricing kernel is defined over realizations of the optimal portfolio of the investor. We provide a framework that unifies the estimation of the pricing kernel and the recovery of the optimal portfolio weights and endogenous wealth. One implication of our method is that any nonmonotonic EPK can be rationalized by a pricing kernel that is monotonically decreasing in the endogenous wealth. Therefore, there is no puzzle when we consider all available investment opportunities.

We show in simulations that the U-shaped EPK arises in the presence of stochastic volatility and jumps risk premia. These can be thought of as additional state variables that make options non-redundant with respect to the index, so that trading in them improves the utility of the investor. In this case, the assumption that the index is perfectly correlated with wealth will be misspecified. A related result from Beare (2011) shows that the cheapest portfolio giving the payoff distribution of the market index is the index itself if and only if the EPK is monotonic. Otherwise, options can be used to construct a portfolio yielding the same payoff distribution as the index at a smaller cost.

Another implication of our framework is that we can interpret the market return states where the EPK is higher as the ones where the endogenous wealth of the marginal investor in the index and index options is smaller. In the case of a U-shaped EPK, this indicates that the marginal investor sells protection against large movements in the index, as extreme negative and positive market returns are associated to smaller endogenous wealth. This is consistent with the model of heterogeneous beliefs about market return outcomes in Bakshi and Madan (2008). While optimistic investors long in the market buy OTM puts as an insurance against a sharp decline in the index, pessimistic investors shorting the index buy OTM calls as an insurance against a rising market. In this sense, the marginal investor in the index and index option markets can be seen as an intermediary satisfying the other side of these positions and harvesting the premium that the individual investors are willing to pay. In fact, Bakshi, Madan and Panayotov (2010) show that a U-shaped EPK is associated to negative expected returns of OTM puts and OTM calls, as they have low expected payout accompanied by high price. An interesting avenue of future research would be to investigate to what extent demand-pressure effects, as in Gârleanu, Pedersen and Poteshman (2009), can explain the time variation in the shape of the EPK and in the optimal portfolio of the marginal investor in the index and index option markets.

3.6 Figures

Figure 3.1: **State-Price Density and Empirical Pricing Kernel - Black-Scholes Economy**

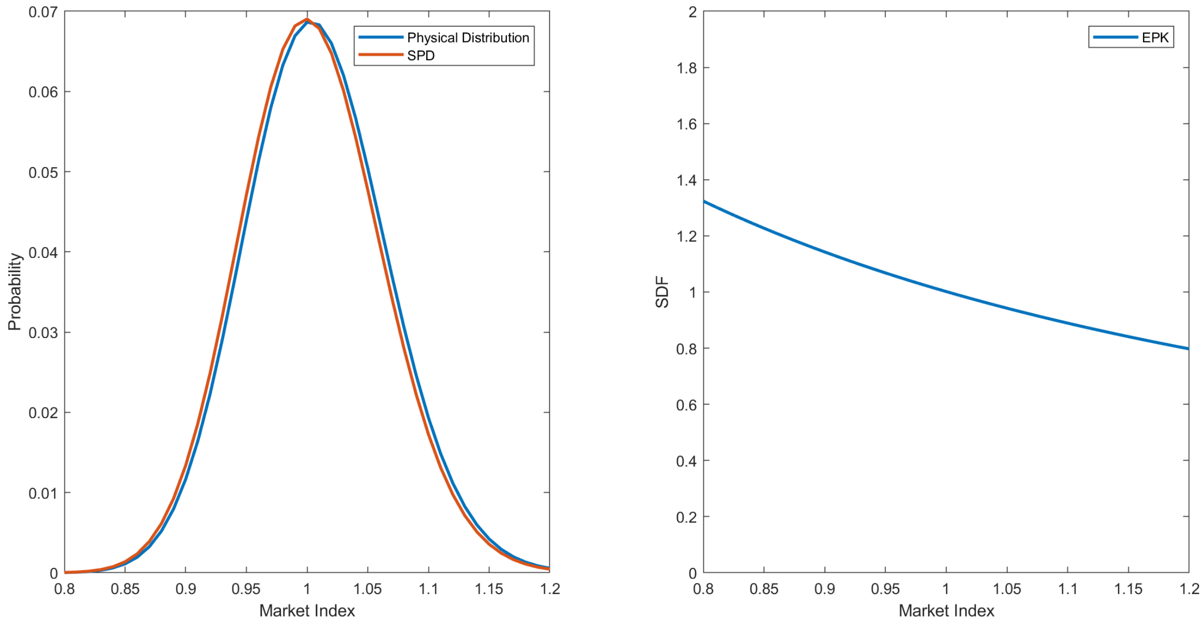

This figure plots for the Black-Scholes economy the risk-neutral and physical distributions, in the left panel, and the empirical pricing kernel, in the right panel. We calculate the risk-neutral and physical distributions as the lognormal densities with the corresponding risk-neutral and physical parameters. The empirical pricing kernel is given by the ratio of the risk-neutral and physical distributions.

Figure 3.2: **Estimated SPD and EPK Using Market Returns - Black-Scholes Economy**

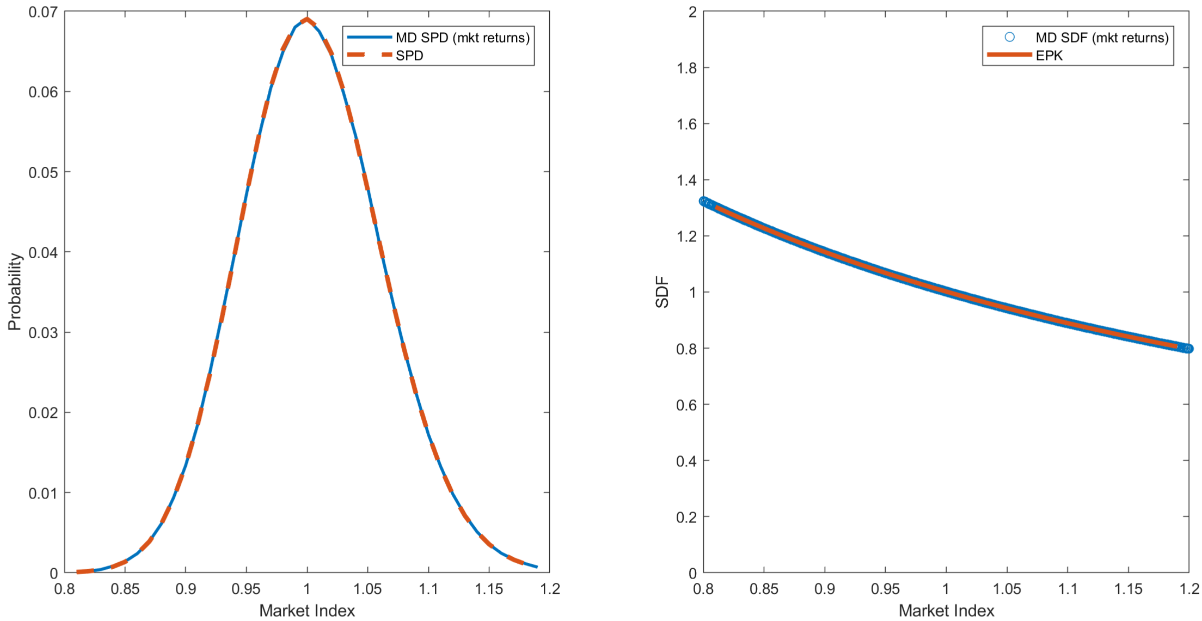

The left panel of this figure plots the SPD coming from the minimum dispersion risk-neutral measure with $\gamma^* = -0.8$ using only market returns (MD SPD (mkt returns)). In red dashed line is the model-implied SPD. The right panel of this figure depicts the model-implied EPK in red and the estimated minimum dispersion SDF projected onto market returns in blue circles.

Figure 3.3: **Estimated Pricing Kernel Using Option Prices - Black-Scholes Economy**

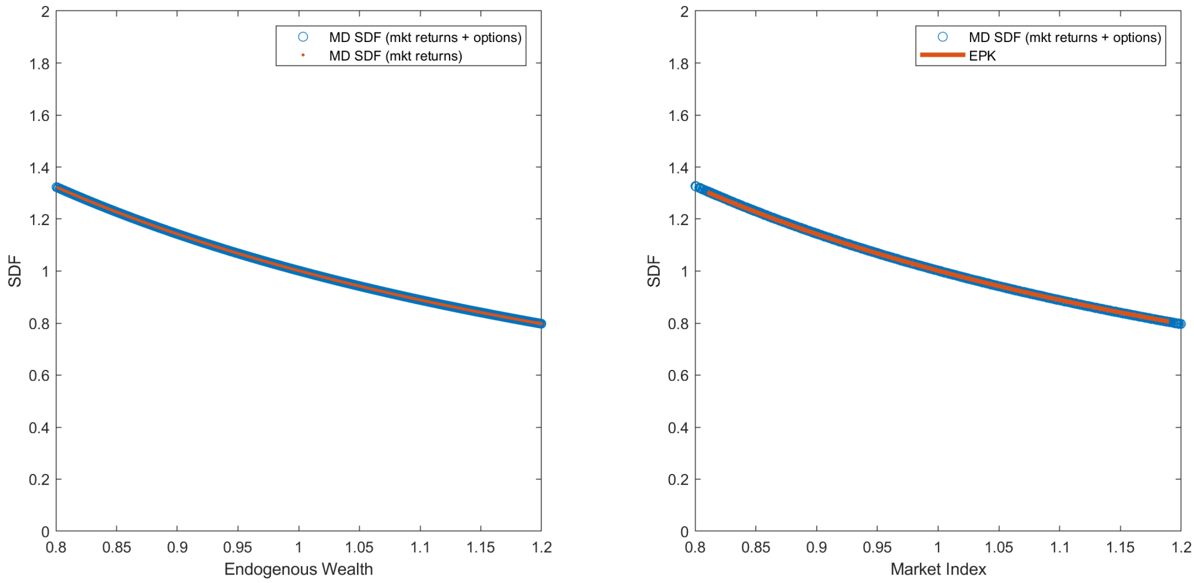

The left panel of this figure plots the minimum dispersion SDF (MD SDF) estimated using market returns and option prices and the MD SDF estimated using only market returns over the endogenous wealth. The right panel plots the MD SDF estimated using market returns and option prices projected onto market returns and the model-implied EPK.

Figure 3.4: **State-Price Density and Empirical Pricing Kernel - SVCJ Economy**

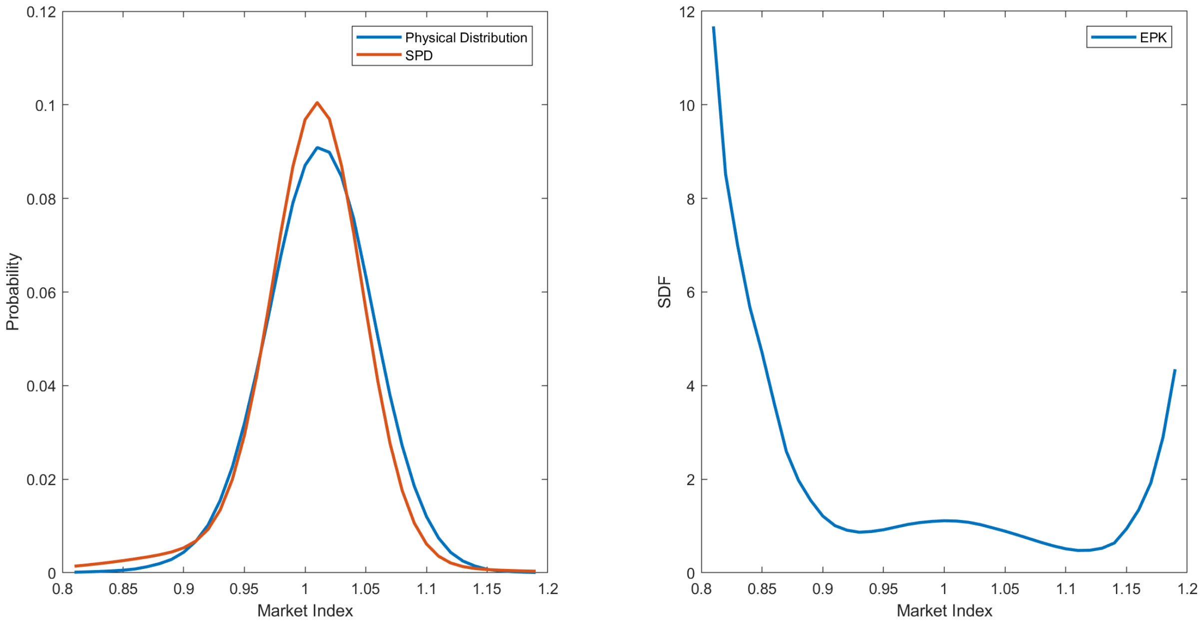

This figure plots for the SVCJ economy the risk-neutral and physical distributions, in the left panel, and the empirical pricing kernel, in the right panel. We calculate the risk-neutral distribution via Breeden and Litzenberger (1978) and the physical distribution deriving the kernel density from sampled underlying returns. The empirical pricing kernel is given by the ratio of the risk-neutral and physical distributions.

Figure 3.5: **Estimated SPD and EPK Using Market Returns - SVCJ Economy**

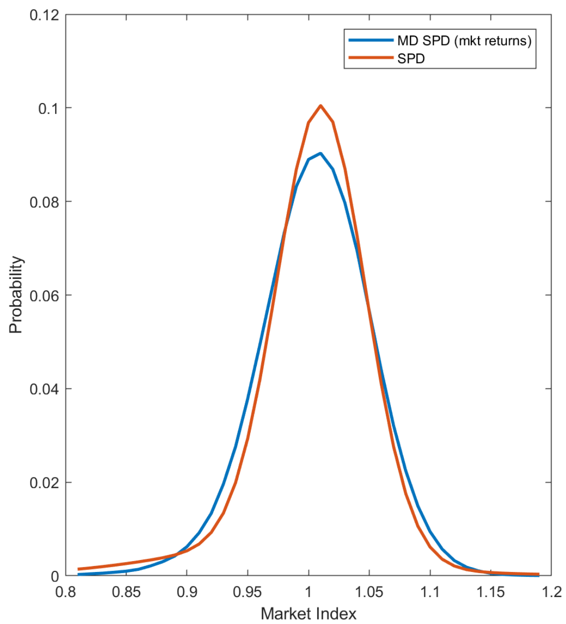

The left panel of this figure plots the SPD coming from the minimum dispersion risk-neutral measure with $\gamma = -1$ using only market returns (MD SPD (mkt returns)). In red is the model-implied SPD. The right panel of this figure depicts the model-implied EPK in red and the estimated minimum dispersion SDF projected onto market returns in blue circles.

Figure 3.6: **Estimated SPD and Pricing Kernel Using Option Prices - SVCJ Economy**

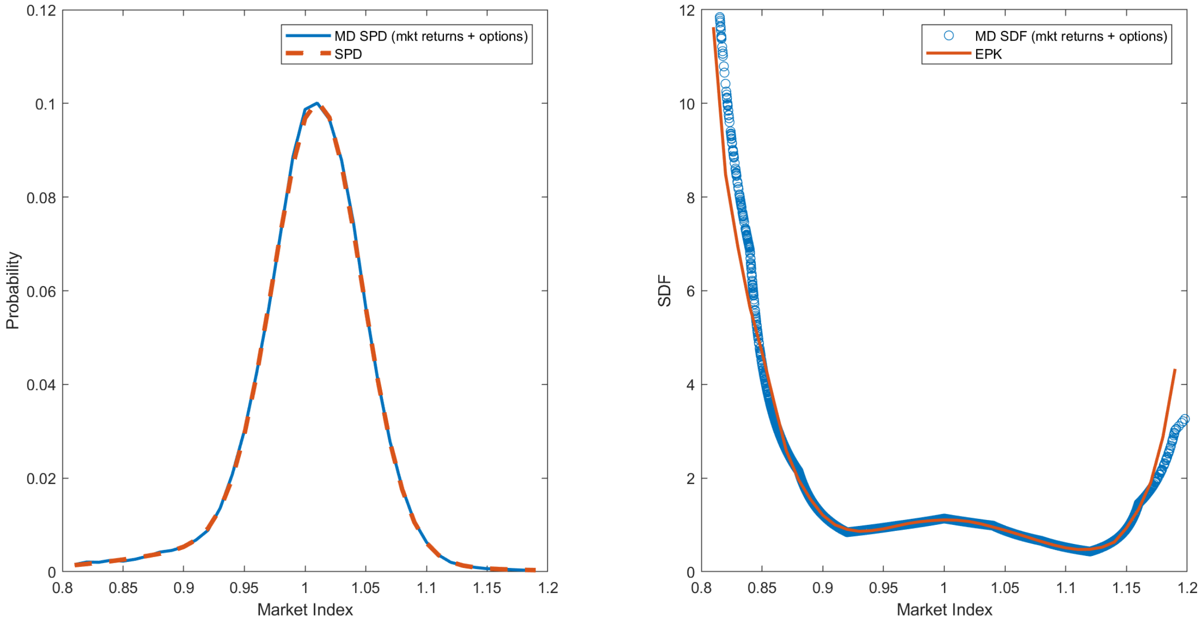

The left panel of this figure plots the SPD coming from the minimum dispersion risk-neutral measure with $\gamma = -1$ using market returns and option prices (MD SPD (mkt returns + options)). In red dashed line is the model-implied SPD. The right panel of this figure depicts the model-implied EPK in red and the estimated minimum dispersion SDF projected onto market returns in blue circles.

Figure 3.7: **Estimated Pricing Kernel Over Endogenous Wealth - SVCJ Economy**

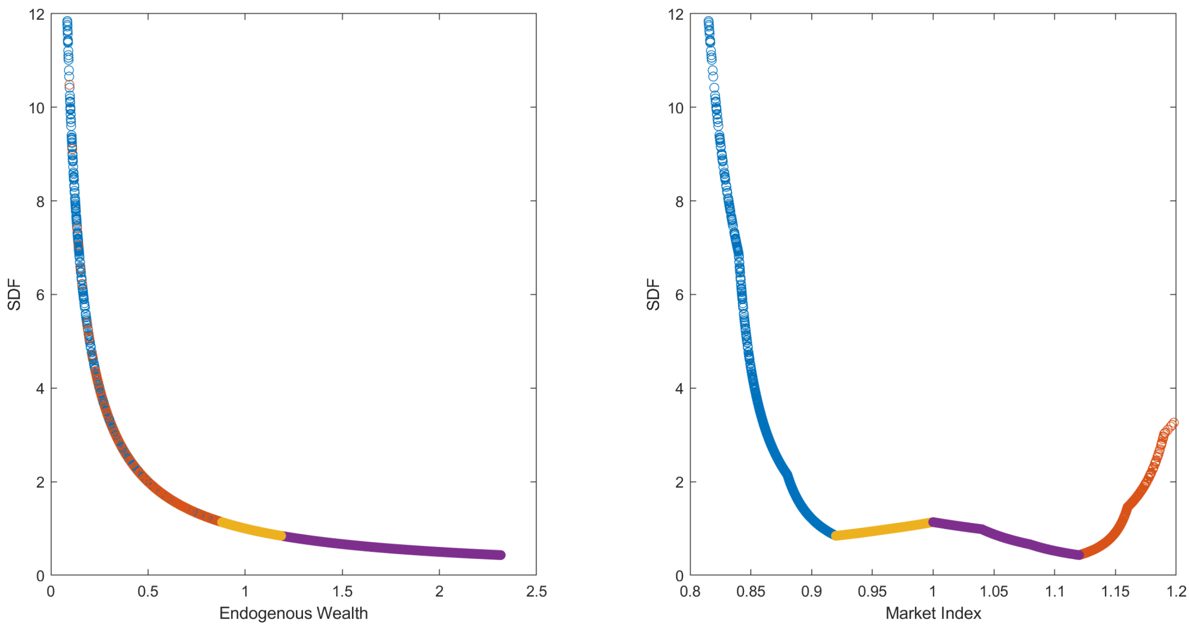

The left panel of this figure plots the estimated pricing kernel over realizations of the endogenous wealth. The right panel plots the same estimated pricing kernel projected onto market returns. Each color represents the same set of realizations of states of nature.

Chapter 4

Tail Risk and Investors' Concerns

In this paper, I estimate tail risk for Brazil and investigate the origins of tail risk variation. The tail risk measure peaks at stock market crashes, financial crises, political shocks and disaster events such as the coronavirus pandemic. Moreover, I find that tail risk is countercyclical, has strong predictive power for market returns and negatively predicts real economic activity. In order to identify the investors' concerns associated to tail risk, I extract daily news from the largest financial newspaper in Brazil. The co-movement between news and tail risk indicates that tail risk variation is mainly driven by disaster concerns, followed by economic and government uncertainty. While economic uncertainty explains the countercyclical property of tail risk, investors only require compensation for bearing tail risk implied by disaster concerns. Similarly, tail risk negatively impacts real outcomes because of the disaster concerns that it identifies. These findings support recent models explaining asset pricing puzzles with time-varying disaster risk.

4.1 Introduction

Since Rietz (1988), a number of papers have shown that left-tail shocks and disaster risk help explain asset pricing behavior that was puzzling from the perspective of traditional macro-finance models (Barro, 2006; Gabaix, 2012; Gourio, 2012; Wachter, 2013). Closely related, there is a literature proposing methods to estimate time-varying tail risk in order to investigate its effects on the stock market and the real economy (Bollerslev and Todorov, 2011; Kelly and Jiang, 2014; Almeida et al., 2017; Weller, 2019). Evidence from the U.S. indicates that tail risk significantly relates to data in the manner predicted by disaster risk models.

The literature has in common that tail risk is estimated based on the idea that asset prices reflect investors' concerns about future states of the economy.[1] However, it is not clear what types of concerns drive tail risk fluctuations. In other words, the origins of tail risk variation are mostly unknown. Identifying such concerns would be important to map how they relate to tail risk premia and affect the real economy, besides learning about economic agents' perception of tail risk. It can also provide novel insights on the modeling of disaster risk and help guide economic policy towards minimizing tail risk.[2]

Moreover, while the effects of time-varying tail risk have been documented for developed economies, the evidence for emerging markets is scarce. Arguably, emerging economies are subject to more uncertainty shocks, representing natural environments to study tail risk. This lack of evidence may be partially explained by the fact that most available methods to measure extreme event risk are data intensive. Since emerging markets are relatively less liquid than their developed counterparts, this usually poses a challenge to construct viable measures of tail risk.

In this paper, I contribute to fill the gaps discussed above. First, I estimate tail risk for Brazil, an important emerging economy that in the past decades has gone through a series of events potentially related to tail risk. Examples include economic recessions, a presidential impeachment, corruption scandals and the coronavirus pandemic. I extract tail risk from a cross-section of portfolio returns following the method in Almeida et al. (2017), which is especially suited for illiquid markets. The estimated monthly tail risk measure peaks at stock market crashes, financial crises, episodes related to government

[1] Bollerslev and Todorov (2011) use intraday futures data and the cross-section of S&P 500 options to construct an investors' fear index. Kelly and Jiang (2014) estimate tail risk from the cross-section of individual stock returns. Almeida et al. (2017) introduce a tail risk measure based on the risk-neutral excess expected shortfall of a cross-section of portfolio returns. Weller (2019) exploits information in the cross-section of bid-ask spreads to develop a measure of extreme event risk.

[2] For instance, minimizing concerns about tail risk is present in the speech of Ben Bernanke at the Federal Reserve Bank of Kansas City Economic Symposium on August 31, 2012: "Such signaling [of large-scale asset purchases] can also increase household and business confidence by helping to diminish concerns about 'tail' risks."

uncertainty and at the coronavirus pandemic. I find that tail risk is countercyclical and has strong predictive power for aggregate market returns, both in-sample and out-of-sample, even after controlling for standard risk factors. Furthermore, employment and industrial production drop significantly following a shock in tail risk, even after controlling for a volatility shock as in Bloom (2009). These results indicate that tail risk strongly affects the stock market and economic activity in Brazil.

Second, I investigate the origins of tail risk fluctuations. Following Manela and Moreira (2017), I consider the words chosen by the business press as a proxy for investors' concerns.[3] I collect daily news from the largest financial newspaper in Brazil, *Valor Econômico*, in order to estimate the co-movement between monthly frequencies of 21,460 words and tail risk. This high-dimensional regression is estimated using elastic net (Zou and Hastie, 2005), a regularization technique from machine learning. I find that tail risk is associated to heterogeneous investors' concerns. In particular, I group words in three different categories, representing concerns related to disasters, economic and government uncertainty. Most of the variation in tail risk is driven by disaster concerns (61.55%), followed by economic (16.32%) and government (6.43%) uncertainty.[4]

The decomposition of tail risk implied by each category of words allows for the identification of the investors' concerns that are responsible for the strong effects of tail risk on aggregate market returns and real outcomes. In particular, economic uncertainty explains the countercyclical property of tail risk, while tail risk premia is mainly associated to disaster concerns. That is, investors only require compensation for bearing tail risk implied by uncertainty about disasters. Similarly, tail risk negatively impacts employment and industrial production because of the disaster concerns that it identifies. Overall, these findings support a recent literature incorporating time-varying disaster risk to solve a number of asset pricing puzzles (Gabaix, 2012; Gourio, 2012; Wachter, 2013).

This paper is mainly related to three strands of the literature. First, to the already cited literature that investigates the effects of tail risk on asset markets and economic activity. I contribute with new evidence from an emerging economy with a relatively illiquid market. The results suggest that investors are tail risk averse, requiring higher returns to hold the market when tail risk increases, which induces a positive relation between tail risk and future market returns. Furthermore, a tail risk shock increases uncertainty, which negatively impacts real outcomes. In line with Bloom (2009), because of capital and labor adjustment costs, this uncertainty shock increases firms' real-option value of postponing investment and hiring decisions. This results in decreased aggregate employment and output.

[3]This is consistent with the model of media bias in Gentzkow and Shapiro (2006) and with the idea that news media reflect the interests of readers, as in Tetlock (2007) and Manela (2011).

[4]The remaining variation (15.67%) is related to an unclassified category.

The second literature is one that extracts information about aggregate uncertainty from news coverage. The word-list based measure of Loughran and McDonald (2011) and the economic policy uncertainty index of Baker, Bloom and Davis (2016) use a human-centric approach to extract this information. Using machine learning, Manela and Moreira (2017) construct a text-based measure of uncertainty that predicts the VIX index from newspaper articles backwards in time. Similarly, I use machine learning to extract the tail risk implied by words chosen by the business press. I obtain new results on how tail risk premia and the impacts of extreme event risk on economic activity can be traced back to heterogeneous investors' concerns. More specifically, concerns related to disasters are of the utmost importance.

This paper is also related to a growing literature applying machine learning methods to overcome the high dimensionality that naturally arises in financial economics problems. To cite a few, Kozak, Nagel and Santosh (2020) use elastic net to estimate an SDF from a large cross-section of returns, Gu, Kelly and Xiu (2020) apply a variety of machine learning techniques to predict stock returns, and Chen, Pelger and Zhu (2020) estimate an asset pricing model with deep neural networks. In this paper, I show how elastic net can be used to estimate the co-movement between news coverage and tail risk. More than that, my approach highlights how the interpretability of text-based features can be useful.

The remaining of the paper is organized as follows. Section 4.2 describes the methodology adopted for constructing a tail risk measure and estimating its relation with topics covered by the business press. Section 4.3 describes the data used throughout the paper. Section 4.4 presents and discusses the empirical results. Section 4.5 concludes the paper.

4.2 Methodology

In this section, I present the methodology adopted in the paper. First, I describe how to estimate a minimum discrepancy risk-neutral measure from a set of basis assets returns and use it to calculate tail risk, following Almeida et al. (2017). Then, in order to identify the investors' concerns driving tail risk fluctuations, I discuss how to estimate the co-movement between news coverage and tail risk.

4.2.1 Minimum Discrepancy Risk-Neutral Measure

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and $\mathbf{R}$ denote a K-dimensional random vector on this space representing the gross returns of a set of K basis assets. Denote by $\mathbf{R}^e \equiv$

$\mathbf{R} - \mathbf{1}_K$ the vector of excess returns, where $\mathbf{1}_K$ is a conformable vector of ones.[5] In this setting, an admissible stochastic discount factor (SDF) is a random variable m satisfying the Euler equations:

$$\mathbb{E}(m\mathbf{R}^e) \equiv \int m\mathbf{R}^e \, \mathrm{d}\mathbb{P} = \mathbf{0}_K, \tag{4.1}$$

where $\mathbf{0}_K$ is a conformable vector of zeros and $\mathbb{P}$ is the physical probability measure. There is a direct correspondence between nonnegative SDFs and risk-neutral measures. Dividing (4.1) by $\mathbb{E}(m)$ and considering the change of measure $\mathrm{d}\mathbb{Q} = \frac{m}{\mathbb{E}(m)} \mathrm{d}\mathbb{P}$, we get the following conditions:

$$\int \mathbf{R}^e \frac{m}{\mathbb{E}(m)} \, \mathrm{d}\mathbb{P} = \int \mathbf{R}^e \, \mathrm{d}\mathbb{Q} \equiv \mathbb{E}^{\mathbb{Q}}[\mathbf{R}^e] = \mathbf{0}_K, \tag{4.2}$$

where the risk-neutral measure $\mathbb{Q}$ is characterized by the state-price density $\mathrm{d}\mathbb{Q}$. An SDF only produces a risk-neutral measure if it is nonnegative. Therefore, the set of admissible risk-neutral measures will be determined by the set of admissible nonnegative SDFs. In complete markets, the absence of arbitrage implies that a unique strictly positive admissible SDF exists. However, in the usual case of incomplete markets, where the number of states of nature in Ω is larger than the number of basis assets K, no-arbitrage implies an infinity of admissible strictly positive SDFs, and, hence, of admissible risk-neutral measures.[6]

In order to nonparametrically estimate an admissible risk-neutral measure from the set of basis assets returns, Almeida et al. (2017) suggest to consider a measure associated to a minimum discrepancy SDF. This approach builds on the seminal work of Hansen and Jagannathan (1991), who propose to obtain an admissible linear SDF by minimizing a quadratic loss function. Almeida and Garcia (2017) generalize this methodology by estimating nonlinear SDFs minimizing a family of discrepancy loss functions (Cressie and Read, 1984). The so-called minimum discrepancy SDFs are especially appropriate for the estimation of a tail risk measure. First, they are nonnegative, guaranteeing the existence of the corresponding risk-neutral measure. Second, they incorporate information about higher-order moments of the basis returns beyond the variance, which is essential to modeling tail risk, since investors' concerns about downside risk are related to negative skewness aversion.[7]

Considering the sample version of the minimum discrepancy problem in Almeida and Garcia (2017), the sample space Ω is discrete and finite, with states of nature $k = \{1, ..., n\}$, where $n > K$. Let $\{\mathbf{R}_k\}_{k=1}^{n}$ be the observed gross returns of the K basis assets, where each $\mathbf{R}_k$ is independent and identically distributed according to $\mathbb{P}$. The

[5] I follow Almeida et al. (2017) in fixing the SDF mean to one.
[6] See Cochrane (2001).
[7] See Schneider and Trojani (2015).

unknown physical measure $\mathbb{P}$ can be replaced by the empirical measure $\mathbb{P}_n$ that gives weights $\pi_k = 1/n$ to the realization of each state of nature.[8] This allows us to exchange the expectation $\mathbb{E}$ with its sample counterpart $\frac{1}{n}\sum_{k=1}^{n} \equiv \sum_{k=1}^{n}\pi_k$. For a given $\gamma \in \mathbb{R}$, the minimum discrepancy problem is given by:

$$\begin{aligned} &\underset{\{m_1,\ldots,m_n\}}{min} \ \tfrac{1}{n}\textstyle\sum_{k=1}^{n} \tfrac{m^{\gamma+1}-1}{\gamma(\gamma+1)}, \\ &\text{s.t. } \tfrac{1}{n}\textstyle\sum_{k=1}^{n} m_k \left(\mathbf{R}_k - \mathbf{1}_K\right) = \mathbf{0}_K, \ \tfrac{1}{n}\textstyle\sum_{k=1}^{n} m_k = 1, \ m_k \geq 0 \ \ \forall\, k. \end{aligned} \tag{4.3}$$

The parameter γ indexes the convex loss function in the Cressie and Read (1984) discrepancy family. This family captures as particular cases several loss functions in the literature, such as the Hansen and Jagannathan (1991) quadratic loss function (γ=1) and the Kullback Leibler Information Criterion ($\gamma \to 0$) considered in Stutzer (1995). Under the assumption of no-arbitrage in the observed sample, Almeida and Garcia (2017) show that solving (4.3) is equivalent to solving the simpler dual problem below, for $\gamma < 0$:

$$\lambda_\gamma^* = \underset{\lambda \in \Lambda_\gamma}{arg\ max}\ \frac{1}{n}\sum_{k=1}^{n} -\frac{1}{\gamma+1}\left(1 + \gamma\lambda'\left(\mathbf{R}_k - \mathbf{1}_K\right)\right)^{\left(\frac{\gamma+1}{\gamma}\right)}, \tag{4.4}$$

where $\Lambda_\gamma = \{\lambda \in \mathbb{R}^K : for\ k = 1, ..., n,\ \left(1 + \gamma\lambda'\left(\mathbf{R}_k - \mathbf{1}_K\right)\right) > 0\}$. The minimum discrepancy SDF can then be recovered from the first-order condition of (4.4) with respect to λ, evaluated at λ_γ^*:

$$m_k^*(\gamma, \mathbf{R}) = \left(1 + \gamma\lambda_\gamma^{*\prime}\left(\mathbf{R}_k - \mathbf{1}_K\right)\right)^{\frac{1}{\gamma}}, \ \ k = 1, ..., n. \tag{4.5}$$

For each γ, the solution λ_γ^* of the dual problem leads to a different minimum discrepancy SDF. The corresponding minimum discrepancy risk-neutral measure is obtained by distorting the empirical measure π_k by the SDF:[9]

$$\pi_k^{\mathbb{Q}^*}(\gamma, \mathbf{R}) = \frac{m_k^*(\gamma, \mathbf{R})}{n}, \ \ k = 1, ..., n. \tag{4.6}$$

4.2.2 Tail Risk Measure

Prices of out-of-the-money (OTM) put options reveal information about the left tail of the risk-neutral return distribution. Thus, it would be natural to use option data to estimate tail risk. However, the option-based approach is data-intensive and often unfeasible in illiquid markets. Almeida et al. (2017) circumvent this issue with risk-

[8] This constitutes an optimal nonparametric estimator for $\mathbb{P}$.

[9] In the sample version, we have that $\frac{\pi_k^{\mathbb{Q}}}{\pi_k} = \frac{m_k}{1/n\sum_{k=1}^{n} m_k}$, where $1/n\sum_{k=1}^{n} m_k = 1$. This implies that the risk-neutral measure is $\pi_k^{\mathbb{Q}} = \pi_k m_k = \frac{m_k}{n}$.

neutralization. They propose to estimate the tail risk of asset i at month t from a set of daily returns $\{R_{i,k}\}_{k=1}^{n}$ as the excess expected shortfall under a minimum discrepancy risk-neutral measure:[10]

$$TR_{i,t} = \sum_{k=1}^{n} \pi_k^{\mathbb{Q}^*}(\gamma, \mathbf{R}) \max(VaR_\alpha(R_{i,k}) - R_{i,k}, 0), \tag{4.7}$$

where $\pi_k^{\mathbb{Q}^*}(\gamma, \mathbf{R})$ is calculated from the basis assets daily returns $\{\mathbf{R}_k\}_{k=1}^{n}$ and α is the Value-at-Risk (VaR) threshold under the risk-neutral distribution. While the tail risk measure above requires only data on a set of assets returns, it still preserves the same intuition of the option-based approach. In particular, the expression in (4.7) is the price of a synthetic OTM put for asset i implied by the minimum discrepancy risk-neutral measure, where the strike price is the α^{th} percentile of the risk-neutral distribution and $\max(VaR_\alpha(R_{i,k}) - R_{i,k}, 0)$ is the put payoff in each state of nature $k = 1, ..., n$ (assuming a synthetic initial price of one for the asset).

As the tail risk measure is estimated at a monthly frequency, the states of nature $k = 1, ..., n$ are captured by previous daily returns starting from the last day of each month t. A relatively small number of days gives the tail risk a quicker response to market conditions, given that recent returns will be used to estimate the risk-neutral measure. I follow Almeida et al. (2017) in considering a window of past thirty business days to estimate tail risk for each month. For the VaR threshold, I consider 15% in order to have a reasonable number of returns to extract information from. As for the choice of the minimum discrepancy risk-neutral measure, Almeida et al. (2017) conduct a series of robustness tests and conclude that $\gamma = -0.5$ is the most appropriate one for the estimation of tail risk.[11] Therefore, I also choose this value for the parameter γ. The aggregate tail risk is calculated as the average over the basis assets:

$$TR_t = \frac{1}{K} \sum_{i=1}^{K} TR_{i,t}. \tag{4.8}$$

4.2.3 News Coverage and Tail Risk

A tail risk measure builds on the idea that prices reflect investors' concerns about future states of the economy. However, there is a gap between measuring tail risk and identifying which concerns are associated to tail risk and how. In order to investigate the origins of tail risk fluctuations, I follow Manela and Moreira (2017) in considering the

[10] The expected shortfall is a coherent measure of risk and contains information about the whole tail of the distribution instead of just a point-wise percentile as VaR.

[11] They also rely on the theoretical results of Kitamura, Otsu and Evdokmov (2013), who show that this choice of γ is the most robust based on an asymptotic perturbation criterion.

time variation in the choice of words by the business press as a proxy for the evolution of the concerns of the average investor. This assumption is consistent with the model of media bias in Gentzkow and Shapiro (2006) and with the idea that news media reflect the interests of readers (Tetlock, 2007; Manela, 2011). Therefore, I fit a model that relates news coverage with tail risk in order to identify the investors' concerns representing sources of tail risk variation.

I begin by constructing a news dataset consisting of the abstracts of daily front-page articles in the online version of the largest financial newspaper in Brazil, *Valor Econômico*.[12] This is accomplished by performing web scraping, a technique employed to extract large amounts of data from websites. I download the *HTML* codes for each daily news coverage and extract the news abstracts from August 2011 to July 2020.[13] The text needs to be treated in order to be used in text-based analysis. First, I eliminate digits, special characters and punctuation. Second, I remove stopwords (highly frequent words) and words that appear less than five times in the whole sample. The remaining text is separately broken into one- and two-word n-grams.[14] Third, I employ Part-of-Speech tagging[15] to classify the part of speech of each word using the *nlpnet* library in Python, which performs natural language processing tasks based on neural networks and provides a Part-of-Speech tagger for the Portuguese language.[16]

Having tagged the words, only nouns are left in the sample, totaling 21,460 n-grams in the dataset. In order to get a relatively large body of text for each observation and to match the frequency of the tail risk measure, the number of times each n-gram appears in each day is counted and the counts are aggregated to the monthly frequency. The number of words per day and per article varies, so I normalize n-gram counts by the total number of n-grams that appear each month, generating a $J = 21,460$ vector $\mathbf{X}_t = [X_{t,1}, ..., X_{t,J}]'$ of n-gram frequencies for each month:

$$X_{t,j} = \frac{\text{appearances of n-gram } j \text{ in month } t}{\text{total n-grams in month } t}. \tag{4.9}$$

The relation between news coverage and tail risk is derived from the co-movement between n-gram frequencies and the estimated tail risk measure. That is, n-gram frequen-

[12]The online version of the newspaper is available at https://valor.globo.com/impresso. For a previous date, as for example, March 2 2020, it can be accessed by the link https://valor.globo.com/impresso/20200302.

[13]The articles in the website were only available since August 2011. The website also went into maintenance from May 2019 to August 2019. For these months, there are no news available.

[14]An n-gram is a contiguous sequence of n words from a given sample of text. For example, "pension" and "reform" are one-word n-grams, while "pension reform" is a two-word n-gram.

[15]Process of marking up a word in a text as corresponding to a particular part of speech, based on both its definition and its context.

[16]The *nlpnet* Part-of-Speech tagger is based on Fonseca and Rosa (2013).

cies are used to explain tail risk with a linear regression model:

$$TR_t = w_0 + w'\mathbf{X}_t + \epsilon_t, \qquad t = 1, ..., T, \tag{4.10}$$

where w is a J vector of regression coefficients and w_0 is the intercept. This regression will be estimated monthly for t ranging from August 2011 to July 2020, totaling $T = 104$ observations. Since $J \gg T$, this linear regression model cannot be estimated using the usual ordinary least squares method. To overcome the high dimensionality of the problem, I employ regularization via elastic net, as briefly described below.[17]

The elastic net is a regularized regression that linearly combines two forms of regularization: LASSO (L^1 norm) and ridge (L^2 norm). The elastic net regression minimizes the following objective function:

$$H(w, w_0) = \sum_{t=1}^{T}(TR_t - w_0 - w'\mathbf{X}_t)^2 + \lambda_1 \sum_{j=1}^{J}|w_j| + \lambda_2(w'w). \tag{4.11}$$

The L^1 penalty employs variable selection by setting elements of $\hat{w}$ to zero, generating a sparse model. However, LASSO selects at most T variables and fails to do grouped selection, i.e., it tends to select one variable from a group of correlated variables and ignore the others. The inclusion of the L^2 penalty removes the limitations on the number of selected variables, encourages grouping effect and stabilizes the L^1 regularization path.[18] Therefore, elastic net works by setting weights for irrelevant variables to zero with the L^1 penalty and at the same time shrinking coefficients of correlated regressors towards each other with the L^2 penalty. This is especially suited for the large dataset of n-grams, since it should contain irrelevant words for tail risk and also correlated words that affect tail risk together. The elastic net regression is estimated numerically and the hyper-parameters λ_1 and λ_2 are chosen by 5-fold cross-validation.

The model above together with the text-based feature of the regressors can provide novel insights into the origins of tail risk fluctuations. In particular, the fitted coefficients $\hat{w}$ supply direct evidence of which words chosen by the business press are associated to tail risk and whether this relation is positive or negative, identifying the types of concerns that the average Brazilian investor relates to tail risk. Furthermore, the same methodology can be applied to other uncertainty measures, allowing us to compare the concerns that each measure captures.

[17] In unreported results available upon request, I also use support vector regression, as in Manela and Moreira (2017). The results are similar to using elastic net.

[18] For more details, see Zou and Hastie (2005).

4.3 Data

The data used to estimate the tail risk measure are daily value-weighted returns on three portfolios sorted by size, three portfolios sorted by book-to-market and three portfolios sorted by momentum, available at NEFIN (Center for Research in Financial Economics), from January 2001 to July 2020.[19] They construct the portfolios for year t considering stocks traded in Bovespa (Brazilian stock exchange) that are eligible according to the following criteria: the stock is the most traded stock of the firm; the stock was traded in more than 80% of the days in year $t-1$ with volume greater than R$ 500.000,00 per day; and the stock was initially listed prior to December of year $t-1$. The construction of these portfolios is based on the methodology by Fama and French (2015), designed to produce spreads in stocks characteristics that are relevant for explaining expected returns in the cross-section. Using such portfolio returns to estimate tail risk has several advantages, as they summarize priced information from the cross-section of returns. This reduces noise and avoids the issue of using an unbalanced panel of stocks.

Throughout the paper, I use data on market returns, risk factors and other uncertainty measures. The market returns considered are monthly returns on the market factor from NEFIN, which is defined by the difference between the value-weighted return of the market portfolio using all eligible stocks and the risk-free rate (computed from the 30-day DI Swap). Risk factors from NEFIN constructed along the lines of Fama and French (2015) will be used as controls in contemporaneous and predictive regressions. The Small Minus Big (SMB) factor is the return of a portfolio long on stocks with low market capitalization (small) and short on stocks with high market capitalization (big). This construction is meant to capture the size effect on average returns. The definition for the other risk factors is analogous. The High Minus Low (HML) is the return of a portfolio long on stocks with high book-to-market and short on stocks with low book-to-market (value factor), Winners Minus Losers (WML) is long on stocks with high past returns and short on stocks with low past returns (momentum factor) and Illiquid Minus Liquid (IML) is long on liquid stocks and short on illiquid stocks (liquidity factor).

Brogaard and Detzel (2015) find that economic policy uncertainty (EPU) positively forecasts excess market returns for the U.S., suggesting that this index is an important risk factor. Therefore, it is pertinent to consider EPU as a potentially relevant risk factor for the Brazilian market as well. I obtain data on the EPU index for Brazil, which is constructed following the methods in Baker, Bloom and Davis (2016) and is available at the website http://www.policyuncertainty.com. They use text archives from the Brazilian newspaper *Folha de São Paulo*. For each month, they count the number of articles containing words considered *a priori* to be relevant for economic policy uncertainty. To

[19]Available at the website http://nefin.com.br.

obtain the EPU rate, they scale the raw EPU counts by the number of all articles in the same newspaper and month.

I also consider the IVol-BR index, available from NEFIN. It consists in an option-implied volatility index for the Brazilian stock market, constructed by Astorino et al. (2017) combining standard methodology with adjustments proposed to take into account the relatively low liquidity of options over Ibovespa (Brazilian stock market index). As the VIX index for the U.S., the IVol-BR can be interpreted as an indicator of market fears in the Brazilian economy. The IVol-BR is provided at a daily frequency and is only available from August 2011 until May 2020, so I average the index over each month to generate monthly figures and consider the smaller sample when comparing with the tail risk measure.

For the economic activity analysis, I obtain daily data on the Ibovespa at Yahoo Finance and monthly data on industrial production and employment for Brazil from the Brazilian Institute of Geography and Statistics (IBGE), from March 2002 to July 2020. Aggregate employment is available from March 2012 to July 2020, while before that there are only employment data on metropolitan areas. I proceed by extrapolating the aggregate employment before March 2012 using the monthly percentage changes in employment in metropolitan areas.

4.4 Empirical Results

4.4.1 Tail Risk Fluctuations in Brazil

Figure 4.1 plots the monthly tail risk series for Brazil from January 2001 to July 2020, estimated according to the methodology in section 4.2.2 using daily returns on nine portfolios sorted by size, value and momentum. The Hoddrick-Prescott (HP) trend component and annotations of the most relevant events are also reported. As can be observed, the tail risk measure peaks at stock market crashes, financial crises, episodes related to political uncertainty and at the coronavirus pandemic, providing a comprehensive depiction of the occurrence of left-tail shocks in the Brazilian economy over time.

The highest peak in tail risk is associated to the COVID-19, confirming the pandemic as a disruptive disaster event. The second and third highest peaks come from the most severe global financial crises over the sample: the subprime crisis and the European sovereign debt crisis. The tail risk measure also responds to a large stock market crash in February 2007 due to the announcement of possible regulations in China's stock market. Political shocks also affect considerably the tail risk series. In particular, the tail risk peaks at key moments such as the JBS delation[20], the presidential impeachment of Dilma Rousseff and the beginning of Lula's presidential term. The 9/11 disaster event is also associated with high tail risk for Brazil. Moreover, the periods with highest trend levels of tail risk are in the beginning of 2001, in the years surrounding the global financial crisis of 2008, in the political crisis around Dilma's impeachment and in the coronavirus pandemic.

It is pertinent to compare the tail risk measure with other measures of uncertainty for Brazil. Figure 4.2 plots the tail risk series and the EPU index for Brazil. Even though their correlation is small (25.31%), the measures share some similarities, usually peaking at the same episodes. The biggest difference between them, however, pertains to the relative magnitude of the response to such episodes. In contrast to the tail risk measure, the EPU is much higher during the political crisis in Brazil from 2015 to 2018 (involving a presidential impeachment and corruption scandals) than in other occasions such as the coronavirus pandemic and the financial crises. This difference may be reflecting that the two measures capture different kinds of uncertainty and, therefore, represent distinct investors' concerns.

Figure 4.3 plots the tail risk measure and IVol-BR, for the smaller sample that the latter is available. As can be seen, the measures are quite similar and display a strong correlation (76.67%). Implied volatility indexes such as the IVol-BR are often interpreted as indicators of investors' crash fears, since option prices reflect investors' expectations about future market conditions. The resemblance between the measures suggests that the risk-neutral excess expected shortfall is able to extract from portfolio returns similar information to that contained in option data, which is by nature forward-looking.

4.4.2 Tail Risk and Market Returns

Kelly and Jiang (2014), Bollerslev, Todorov and Xu (2015) and Almeida et al. (2017) document a premium for bearing tail risk, suggesting that tail risk is a priced risk factor. This is consistent with the idea that investors are tail risk averse, so that increases in aggregate tail risk raise the return required by investors to hold the market portfolio. This naturally induces a positive predictive relation between tail risk and future market returns. Moreover, since tail risk is often countercyclical, we should expect it to be negatively related to contemporaneous market returns.

In this subsection, I investigate the relation between tail risk and market returns in Brazil. I start by analyzing the ability of the tail risk measure in explaining excess

[20] Arguably the most relevant episode of the Car Wash operation (*operação Lava Jato*) in Brazil, an ongoing criminal and corruption investigation by the Brazilian Federal Police. The delation by the executives of the JBS company mentioned illegal payments to over 1800 politicians.

market returns with contemporaneous and predictive regressions:

$$r_t = \alpha + \beta TR_t + \delta' \mathbf{C}_t + \epsilon_t, \qquad t = 1, ..., T, \tag{4.12}$$

$$r_{[t+1,t+h]} = \alpha + \beta TR_t + \delta' \mathbf{C}_t + \epsilon_{[t+1,t+h]}, \qquad t = 1, ..., T-h, \tag{4.13}$$

where r_t is the excess return on the market portfolio at month t, $r_{[t+1,t+h]}$ is the excess return on the market portfolio between $t+1$ and $t+h$ and $\mathbf{C}_t$ includes as controls the EPU index and the risk factors SMB, HML, WML and IML. The forecast horizons considered are $h = 1, 3, 6, 9, 12$ and 24. For each horizon h, the t-statistics for the estimated coefficients are calculated using Newey and West (1987) robust standard errors with a lag length equal to h. For easier interpretation, the results are reported for standardized variables.

Table 4.1 reports the results for regressions (4.12) and (4.13). Focusing first on the contemporaneous regression, tail risk is clearly countercyclical, with a negative coefficient that is strongly significant. The risk factors are also related to contemporaneous market returns, while the coefficient for EPU is not statistically significant. As for the forecast regressions, tail risk positively predicts returns, confirming a premium for bearing tail risk. The coefficients are strongly significant for most of the forecast horizons. The size and liquidity factors also drive risk premium, in contrast with the value and momentum factors. The EPU index only predicts market returns for long horizons, especially for 24 months.

I also investigate the out-of-sample (OOS) predictive power of tail risk, following Kelly and Jiang (2014). Using data only through month t (beginning at $t = 120$ for a sufficiently large initial estimation sample), I run univariate predictive regressions of market returns on tail risk. With the estimated coefficient, I forecast the $t+1$ return. The estimation window is then extended by one month in order to estimate a new coefficient and construct an OOS forecast of the market return in the following month. This is repeated until the full sample has been exhausted. Using the forecast errors, I calculate the OOS R^2 as $1 - \sum_t (r_{[t+1,t+h]} - \hat{r}_{[t+1,t+h]|t})^2 / \sum_t (r_{[t+1,t+h]} - \overline{r}_{[t+1,t+h]|t})^2$, where $\hat{r}_{[t+1,t+h]|t}$ is the OOS forecast of the $t+1$ return using only data through t and $\overline{r}_{[t+1,t+h]|t}$ is the historical average market return through t. A negative OOS R^2 indicates that the predictor performs worse than forecasts given by the historical mean.

Table 4.2 reports the OOS R^2 for tail risk, EPU and risk factors, for horizons $h = 1, 3, 6, 9, 12$ and 24. Tail risk is the only predictor with positive OOS R^2 for all horizons, with statistical significance (at the 5% level using Clark and McCracken's (2001) test) for horizons larger than 3 months. The only predictor with comparable performance is SMB, while the remaining ones do not present substantial predictive power out-of-sample. In particular, EPU only performs better than the historical mean for horizons of

3 and 24 months.

In sum, the results from the regressions suggest that tail risk is countercyclical, while being positively and significantly related to future market returns, even after controlling for EPU and standard risk factors. That is, there is a premium for bearing tail risk in the Brazilian market. Furthermore, the EPU index only presents predictive power for a long horizon of 24 months, in contrast with evidence from the U.S. as in Brogaard and Detzel (2015).

4.4.3 Tail Risk and Economic Activity

Tail risk can affect real outcomes via its impact as an uncertainty shock. Bloom (2009) argues that, in the presence of capital and labor adjustment costs, higher uncertainty increases the real-option value of firms of postponing investment and hiring decisions. This generates a rapid slowdown in economic activity, followed by a bounce-back once uncertainty has decreased and firms have addressed their repressed demand for labor and capital. In particular, Bloom (2009) focuses on uncertainty as measured by stock market volatility, showing that a volatility shock strongly affects industrial production and employment. Kelly and Jiang (2014) consider this framework to assess the impact of tail risk on economic activity. They show that, after controlling for the impact of volatility shocks, a tail risk shock precedes a contraction in real outcomes.

In this subsection, I adopt a similar approach to the papers mentioned above to examine the impact of tail risk on economic activity in Brazil. I estimate a monthly vector autoregression (VAR) from March, 2002 to July, 2020. The variables in the estimation order are log(Bovespa stock market index), stock market volatility, tail risk, log(employment) and log(industrial production).[21] Figure 4.4 plots the response of employment and industrial production (and one standard error bands) to a one-standard-deviation shock to volatility and a one-standard-deviation shock to tail risk. Employment displays a fall of -0.4% after 7 months of the volatility shock, with a subsequent recovery at around two years. Tail risk presents an incremental impact relative to volatility, as a tail risk shock induces a rapid decline in employment of -0.2% within 4 months, and then a bounce-back after two years. The impacts on industrial production are more pronounced. Output decreases by -1% within 3 months of a volatility shock, with a subsequent recovery and rebound from 8 months after. Following a shock to tail risk, industrial production displays an immediate drop of -0.8%, followed by a recovery and rebound from 6 months after. The one standard errors bands highlight that all effects described above are statistically significant at the 5% level.

[21] The stock market volatility is estimated as the standard deviation of daily returns on Ibovespa on month t. Employment and industrial production are seasonally adjusted. All variables are Hodrick-Prescott (HP) detrended with $\lambda = 129,600$ in the baseline estimations.

In summary, a positive shock to tail risk induces a large and prolonged contraction in employment and industrial production in Brazil, even after controlling for the impact of volatility shocks.

4.4.4 Origins of Tail Risk Fluctuations

I now investigate the origins of tail risk fluctuations using the news dataset and the methodology described in section 4.2.3. I will be useful to compare the investors' concerns associated to tail risk to those related to other uncertainty measures, namely the EPU index and the IVol-BR implied volatility index. Following Manela and Moreira (2017), a relevant measure of word importance is the fraction of implied tail risk variance that each word drives. I define $\hat{v}_t(j) \equiv X_{t,j}\hat{w}_j$ as the value of tail risk in month t implied only by n-gram $j \in \{1, ..., J\}$. Therefore, the measure of implied tail risk variance driven by a specific n-gram is given by:

$$h(j) \equiv \frac{var(\hat{v}_t(j))}{\sum_{i \in J} var(\hat{v}_t(i))}. \tag{4.14}$$

Table 3 reports $h(j)$ for the top 20 tail risk, EPU and IVol-BR variance driving n-grams estimated by elastic net from August 2011 to July 2020.[22] A clear pattern can be identified: for the three uncertainty measures, the top variance n-gram is "crisis". That is, an increase in the frequency of this word in front-page news is a strong indication of higher uncertainty. In particular, "crisis" is considerably more important to tail risk, explaining 59.40% of its variation, in comparison with 43.49% and 45.98% for EPU and IVol-BR, respectively. This suggests a strong disaster concern related to the uncertainty that tail risk identifies.

Focusing now on Panel A, we can see that tail risk is related to heterogeneous concerns. Words like "crisis", "pandemic" and "epidemic" can be interpreted as concerns related to the occurrence of disasters, while "economy", "companies", "credit", "firms" and "bonds" to uncertainty about the economy. On the other hand, "minister", "government" and "pension" indicate concerns related to government uncertainty. The evidence is similar for IVol-BR, where the main difference is that "pandemic" drives a larger variance share. In contrast, EPU is mainly associated to government uncertainty, as would be expected. In particular, the top variance words for EPU are closely related to the political crisis in Brazil from 2015 to 2018. The n-gram "impeachment" is associated to the presidential impeachment of Dilma Rousseff, while "car wash" and "operation" to the criminal and corruption investigation led by the Car Wash operation. Also important

[22] Naturally, the words comprising the Brazilian news are in Portuguese. To make it easier for a non-Portuguese speaker to follow the analysis, I report in the main results the words translated into English. In Appendix A 4.6, I report the corresponding original words in Portuguese.

is “reform” and “pension”, referring to the pension reform in Brazil which generated a heated debate in the political scenario.

Overall, in contrast to EPU and similarly to IVol-BR, tail risk captures heterogeneous investors’ concerns, with a stronger presence of a disasters uncertainty component. Further insight can be gained with the estimated coefficients $\hat{w}$ in (4.10), which provide information on the sign of the relation between concerns and tail risk. Figure 4.5 plots the top 25 largest and smallest coefficients for tail risk. The appearance of words like “crisis”, “economy”, “companies”, “minister”, “epidemic”, “credit”, “bonds”, “virus”, “ministry”, “shock”, “pension” and “government” are associated to higher tail risk. On the other hand, n-grams such as “revenues”, “construction”, “productivity”, “inflation”, “consumption” and “performance” lead to smaller tail risk when they appear in front-page news.

In order to analyze in further detail the relation between tail risk and investors’ concerns, I group words in categories representing distinct concerns. To that end, I search, between the top n-grams explaining 99% of tail risk variation, the words representing concerns about disasters, economic and government uncertainty.[23] Table 4.4 summarizes the information regarding the word categories.[24] The disaster category explains the most part of tail risk variation (61.55%), mainly because of the word “crisis”. Words related to economic uncertainty explain 16.32%, while government concerns account for 6.43% of the variation. The remaining unclassified words have a variance share of 15.67%.

Figure 4.6 plots the tail risk implied by each category.[25] An interesting pattern is that the disaster implied tail risk is usually smaller than that implied by economic concerns, except in the extreme event of the coronavirus pandemic. In fact, all categories peak at the beginning of COVID-19, but the highest peak comes from the disaster category. In particular, the government concerns imply higher tail risk after the presidential election of 2018 than during the pandemic. Moreover, the unclassified category is usually contributing to decrease tail risk.

The decomposition of tail risk implied by different categories allows for the identification of the investors’ concerns responsible for the strong effects of tail risk on the stock market and the real economy. In the remaining of this section, I investigate in detail which concerns drive these effects.

---

[23] In particular, 165 words explain 99% of the variation in tail risk.

[24] Appendix A 4.6 reports all words comprising the disaster, economic and government category.

[25] For each category c, the implied tail risk is given by $\widehat{TR}_t(c) \equiv \hat{w}(c)'\mathbf{X}_t$, where $\hat{w}(c)$ contains the estimated coefficients for words in category c and is zero elsewhere.

4.4.5 Tail Risk Premia Decomposition

The results in subsection 4.4.2 indicate that tail risk is countercyclical and positively related to future market returns. In this subsection, I investigate the investors' concerns explaining these relations. I consider contemporaneous and predictive regressions of excess market returns on the categories implied tail risk. Table 4.5 reports the results. From Panel A, we can see that the strong negative relation between market returns and tail risk can be explained by the economic concerns that the tail risk measure captures. Besides the economic category, the unclassified category also presents a significant negative coefficient, while the disaster concerns are positively related to market returns. The coefficient for the government category is not statistically significant.

In the predictive regressions of Panel B, besides the estimated coefficients and t-statistics, I also report the share of risk premia variation due to each of the categories.[26] Perhaps surprisingly, investors do not require compensation for tail risk coming from economic uncertainty, for any horizon. In contrast, disaster concerns explain a large portion of risk premia variation and are strongly and positively related to future market returns, for all forecast horizons except 1 month. Government concerns also drive tail risk premium at a 24-month horizon. The unclassified category is important for horizons of 9 and 12 months. In sum, tail risk premia is primarily driven by disaster concerns.

4.4.6 Investors' Concerns and Economic Activity

I now analyze the extent to which the strong effects of tail risk on real economic activity can be explained by the investors' concerns identified by the tail risk measure. I estimate a monthly vector autoregression (VAR) from August, 2011 to July, 2020 for each category implied tail risk. The variables in the estimation order are log(Bovespa stock market index), stock market volatility, category implied tail risk, log(employment) and log(industrial production). Figure 4.7 plots, for each VAR estimation, the response of employment and industrial production (and one standard error bands) to a one-standard-deviation shock to a category implied tail risk. As can be seen, the impacts on real outcomes of tail risk implied by disaster concerns are large, significant and very similar to those of tail risk in Figure 4.4. A positive shock in economic implied tail risk also affects employment, generating a -0.3% decline within 5 months, and a bounce-back one year and a half later. However, its impact on industrial production is much less pronounced. Similarly, the effects of a shock in tail risk implied by government concerns are small and mostly non-significant.

In summary, only the increase in uncertainty about disasters coming from a tail

[26] The share of tail risk premia variation due to each category c is defined by $cov(\beta_c \widehat{TR}_t(c), \sum_c \beta_c \widehat{TR}_t(c))/var(\sum_c \beta_c \widehat{TR}_t(c))$.

risk shock influences the decision of firms of postponing investment and hiring decisions, leading to a contraction in real outcomes. That is, the strong effects of tail risk on real economic activity can be explained by the disaster concerns that it captures.

4.5 Conclusion

I estimate tail risk for Brazil based on the risk-neutral excess expected shortfall of returns on portfolios sorted by size, book-to-market and momentum. This methodology, developed by Almeida et al. (2017), avoids the shortcomings of data-intensive approaches, being particularly useful for a relatively illiquid market. The estimated tail risk measure peaks at stock market crashes, financial crises, episodes related to political uncertainty and disaster events, providing a comprehensive depiction of the occurrence of extreme events in the Brazilian economy over time. Moreover, I show that tail risk is countercyclical, has strong predictive power for future market returns and negatively predicts economic activity.

In order to investigate the origins of tail risk fluctuations, I consider the time variation in topics covered by the business press as a proxy for the evolution of investors' concerns regarding these topics. I extract news from the largest financial newspaper in Brazil and estimate the co-movement between the news coverage and tail risk. I find that tail risk is mainly associated to three different categories or words, representing concerns related to uncertainty about disasters, the economy and the government. Decomposing tail risk implied by each category, it is possible to identify the investors' concerns responsible for the strong effects of tail risk on the stock market and real economic activity. While economic uncertainty explains the countercyclical property of tail risk, investors only require compensation for bearing tail risk implied by disaster concerns. Similarly, only the increase in uncertainty about disasters coming from a tail risk shock impacts employment and industrial production. These findings support recent models incorporating time-varying disaster risk to explain asset pricing behavior.

This paper provides a number of implications for future research. First, the empirical results indicate that the estimated tail risk measure should be considered as a benchmark risk factor for the Brazilian economy. This is in contrast to the EPU index, which does not appear to be significantly related to aggregate market returns. Second, the identification of the investors' concerns and sources of uncertainty that drive tail risk premia and affect the real economy provide important insights to the disaster risk literature. In particular, it is suggestive of the types of extreme events that are important to consider in calibrating the models. Third, while I focus in interpreting the co-movement between news coverage and tail risk and investigating its implications, a next step would be to assess whether tail risk can be predicted out-of-sample by the words chosen by the business press. This can be particularly useful for economic policy. Fourth, it would be interesting to investigate the origins of high frequency variation in tail risk and compare with the results in this paper. More specifically, a daily tail risk measure can be estimated from intra-day returns using the same methodology (see Almeida et al., 2020). In this case, news data from a cross-section of newspapers would be ideal in order to have a large body of text for each day and estimate the relation with tail risk.

4.6 Appendix A - Word Categories and Original Words

In this appendix, I begin by reporting the words comprising the disaster, economic and government categories, in decreasing order of importance (as measured by implied tail risk variance share).

i) Disaster category: crisis, pandemic, epidemic, quarantine, virus, shock.

ii) Economic category: economy, companies, credit, firms, sector, bonds, inflation, prices, price, production, revenues, costs, fees, fed, consumption, performance, stock exchange, energy, productivity, infrastructure, slowdown, sectors, gdp, investment, products, technology, levy, economists, reserves, factories, exchange rate, monetary committee, employment, industry, financing, constructions, exports, budget, construction, stocks, services, exchanges, stock.

iii) Government category: minister, government, pension, ministry, governors, candidate, round, election, justice, reforms, law, governments, constitutional amendment, public ministry, trial, pension reform, attorney, investigation, federal court.

I now present the original words, in Portuguese, for all translated words appearing in the paper. In alphabetic order, I report the translated word and then the original word in parenthesis:

agency (órgão), attorney (procurador), bonds (títulos), budget (orçamento), candidate (candidato), car wash (lava jato), chamber (câmara), companies (empresas), constitutional amendment (pec), construction (obra), constructions (obras), consumption (consumo), costs (custos), countries (países), credit (crédito), crisis (crise), day (dia), days (dias), december (dezembro), economists (economistas), economy (economia), election (eleição), employment (emprego), energy (energia), epidemic (epidemia), exchange rate (câmbio), exchanges (bolsas), exports (exportações), factories (fábricas), february (fevereiro), federal court (stf), fees (taxas), fed (fed), financing (financiamento), firms (companhias), gdp (pib), government (governo), governments (governos), governors (governadores), group (grupo), hand (mão), impeachment (impeachment), importance (importância), industry (indústria), inflation (inflação), infrastructure (infraestrutura), investigation (investigação), investment (investimento), july (julho), june (junho), justice (justiça), law (lei), levy (levy), lines (linhas), march (março), mark (marco), may (maio), millions (milhões), minister (ministro), ministry (ministério), moment (momento), monetary committee (copom), network (rede), operation (operação), pandemic (pandemia), payment (pagamento), pension (previdência), pension reform (reforma previdência), people (pessoas), performance (desempenho), quarantine (quarentena), period (período), price (preço), prices (preços), process (processo), production (produção), products (produtos), productivity (produtividade), project (projeto), public ministry (mp), quarter (trimestre), raise (aumento), recovery (recuperação), reduction (redução), reform (reforma), reforms (reformas), region (região), reserves (reservas), revenues (faturamento), round (turno), sector (setor), sectors (setores), september (setembro), services (serviços), slowdown (desaceleração), shock (choque), solution (solução), stock (ação), stock exchange (bolsa), stocks (ações), study (estudo), technology (tecnologia), trial (julgamento), value (valor), virus (vírus), year (ano), years (anos).

4.7 Tables

Table 4.1: **Market Returns Regressions**

	Tail Risk	EPU	SMB	HML	WML	IML	adj. R^2
	Panel A: Contemporaneous						
	-0.33***	0.02	0.41***	0.11*	-0.16**	-0.39***	0.26
	(-5.69)	(0.41)	(4.18)	(1.86)	(-2.51)	(-4.11)	
	Panel B: Forecasts						
1 month	0.10**	0.04	0.39***	-0.03	-0.03	-0.27**	0.06
	(2.07)	(0.73)	(2.89)	(-0.49)	(-0.53)	(-2.51)	
3 months	0.20***	0.08	0.39***	-0.03	-0.02	-0.24**	0.08
	(2.96)	(1.04)	(3.73)	(-0.51)	(-0.38)	(-2.25)	
6 months	0.14	0.12	0.36***	-0.06	-0.04	-0.18**	0.06
	(1.47)	(1.47)	(4.19)	(-1.13)	(-0.67)	(-2.14)	
9 months	0.24**	0.13	0.37***	-0.06	-0.01	-0.22**	0.08
	(2.27)	(1.48)	(3.90)	(-1.20)	(-0.26)	(-2.32)	
12 months	0.14	0.16*	0.31***	-0.06	-0.05	-0.13*	0.07
	(1.22)	(1.74)	(3.42)	(-1.31)	(-1.15)	(-1.90)	
24 months	0.20**	0.22**	0.18*	0.03	0.01	-0.07	0.07
	(1.95)	(2.06)	(1.72)	(0.80)	(0.26)	(-1.00)	

This table presents the results for multivariate contemporaneous and forecasting regressions for the excess market return. Panel A reports, for the contemporaneous regression, the coefficients and the t-statistics in parenthesis for the regressors indicated by the columns. The last column reports the adjusted R^2. Panel B reports, for forecasting regressions with horizon indicated by the row, the coefficients and the t-statistics in parenthesis for the regressors indicated by the columns. The t-statistics are calculated using Newey and West (1987) robust standard errors with a lag length equal to the forecast horizon. The last column reports the adjusted R^2. The coefficients for the intercept are omitted. The sample ranges from January, 2001 to July, 2020. *, ** and *** represent statistical significance at the 10%, 5% and 1% levels, respectively.

Table 4.2: **Market Return Predictability: Out-of-Sample R^2 (%)**

	Tail Risk	EPU	SMB	HML	WML	IML
1 month	0.2	-0.7	0.9*	-2.9	2.4*	-0.4
3 months	5.3*	1.1*	-3.8	-0.8	0.6	-1.3
6 months	1.5*	-0.5	4.1*	-1.8	0.9*	0.4
9 months	8.3*	-2.8	2.6*	-2.8	-1.5	-1.2
12 months	2.4*	-4.1	6.2*	-4.6	1.7*	1.3*
24 months	7.5*	1.4*	1.2	-0.6	-0.3	-0.6

This table reports the out-of-sample (OOS) R^2 in percent from predictive regressions of market returns over horizons indicated by the rows. In each month t, beginning at $t = 120$, I estimate rolling univariate forecasting regressions of market returns on each regressor indicated by the columns. The estimated coefficients are used to forecast returns at $t + 1$. The OOS R^2 is calculated as $1 - \sum_t (r_{[t+1,t+h]} - \hat{r}_{[t+1,t+h]|t})^2 / \sum_t (r_{[t+1,t+h]} - \overline{r}_{[t+1,t+h]|t})^2$. The sample ranges from January, 2001 to July, 2020. * indicates statistical significance at the 5% level based on the Clark and McCracken (2001) ENC-NEW test of OOS predictability.

Table 4.3: **Top Variance n-grams for Tail Risk, EPU and IVol-BR**

Panel A: Tail Risk			
n-gram	Variance Share (%)	n-gram	Variance Share (%)
crisis	59.40	credit	0.86
economy	6.02	epidemic	0.86
companies	5.57	march	0.79
minister	3.08	sector	0.63
millions	2.30	people	0.59
years	2.06	firms	0.57
may	1.76	raise	0.52
government	1.30	bonds	0.52
pandemic	1.08	days	0.51
pension	0.97	quarter	0.45

Panel B: EPU			
n-gram	Variance Share (%)	n-gram	Variance Share (%)
crisis	43.49	production	1.03
impeachment	9.24	countries	0.93
car wash	7.46	operation	0.91
millions	4.58	may	0.82
reform	4.23	ministry	0.80
recovery	3.54	pension	0.64
government	2.77	value	0.62
february	1.67	process	0.61
chamber	1.23	group	0.50
reforms	1.13	region	0.50

Panel C: IVol-BR			
n-gram	Variance Share (%)	n-gram	Variance Share (%)
crisis	45.98	reform	0.42
pandemic	28.14	may	0.41
year	5.30	june	0.41
companies	4.33	prices	0.35
economy	2.52	minister	0.32
impeachment	1.83	days	0.30
millions	0.87	september	0.27
round	0.79	firms	0.26
quarter	0.68	bonds	0.21
group	0.66	candidate	0.19

This table reports the share of implied tail risk, EPU and IVol-BR variance that each n-gram drives, for the top 20 n-grams. The sample period is August, 2011 to July, 2020.

Table 4.4: **Word Categories**

Category	Variance Share (%)	n-grams	Top n-grams
Disaster	61.55	6	crisis, pandemic, epidemic, quarantine
Economic	16.32	43	economy, companies, credit, firms
Government	6.43	19	minister, government, pension, ministry
Unclassified	15.67	21,392	millions, years, may, march

This table reports the share of implied tail risk variance that each n-gram category drives ($h_c(j) = \sum_{j \in c} h(j)$), the number of n-grams and the top n-grams in each category.

Table 4.5: **Tail Risk Premia Decomposition**

	Disaster	Economic	Government	Unclassified	adj. R^2
			Panel A: Contemporaneous		
	0.29**	-0.35***	0.07	-0.32***	0.14
	(2.39)	(-2.99)	(0.72)	(-2.68)	
			Panel B: Forecasts		
1 month	0.14	0.08	0.10	-0.02	0.02
	(1.16)	(0.63)	(1.39)	(-0.20)	
	[0.53]	[0.25]	[0.27]	[-0.05]	
3 months	0.30**	0.01	0.08	0.12	0.13
	(2.04)	(0.01)	(0.86)	(0.81)	
	[0.63]	[0.00]	[0.12]	[0.24]	
6 months	0.48**	-0.13	0.12	0.28	0.10
	(2.16)	(-0.72)	(1.10)	(1.32)	
	[0.44]	[0.06]	[0.18]	[0.30]	
9 months	0.51***	-0.25	0.10	0.49**	0.21
	(3.61)	(-1.17)	(1.12)	(2.59)	
	[0.25]	[0.15]	[0.10]	[0.47]	
12 months	0.47***	-0.27	0.09	0.40**	0.16
	(2.73)	(-1.03)	(0.73)	(2.15)	
	[0.27]	[0.21]	[0.10]	[0.40]	
24 months	0.60***	-0.18	0.34**	0.44	0.36
	(2.64)	(-0.95)	(2.41)	(1.71)	
	[0.32]	[0.05]	[0.31]	[0.31]	

This table presents the results for multivariate contemporaneous and forecasting regressions for the excess market return, where the regressors are categories implied tail risk. Panel A reports, for the contemporaneous regression, the coefficients and the t-statistics in parenthesis for the regressors indicated by the columns. The last column reports the adjusted R^2. Panel B reports, for forecasting regressions with horizon indicated by the row, the coefficients, the t-statistics in parenthesis and the variance share in brackets for the regressors indicated by the columns. The t-statistics are calculated using Newey and West (1987) robust standard errors with a lag length equal to the forecast horizon. The variance share is the percent of risk premia variation due to each category, i.e., $cov(\beta_c \widehat{TR}_t(c), \sum_c \beta_c \widehat{TR}_t(c))/var(\sum_c \beta_c \widehat{TR}_t(c))$. The last column reports the adjusted R^2. The coefficients for the intercept are omitted. The sample ranges from August, 2011 to July, 2020. *, ** and *** represent statistical significance at the 10%, 5% and 1% levels, respectively.

4.8 Figures

Figure 4.1: **Tail Risk Measure for Brazil (2001-2020)**

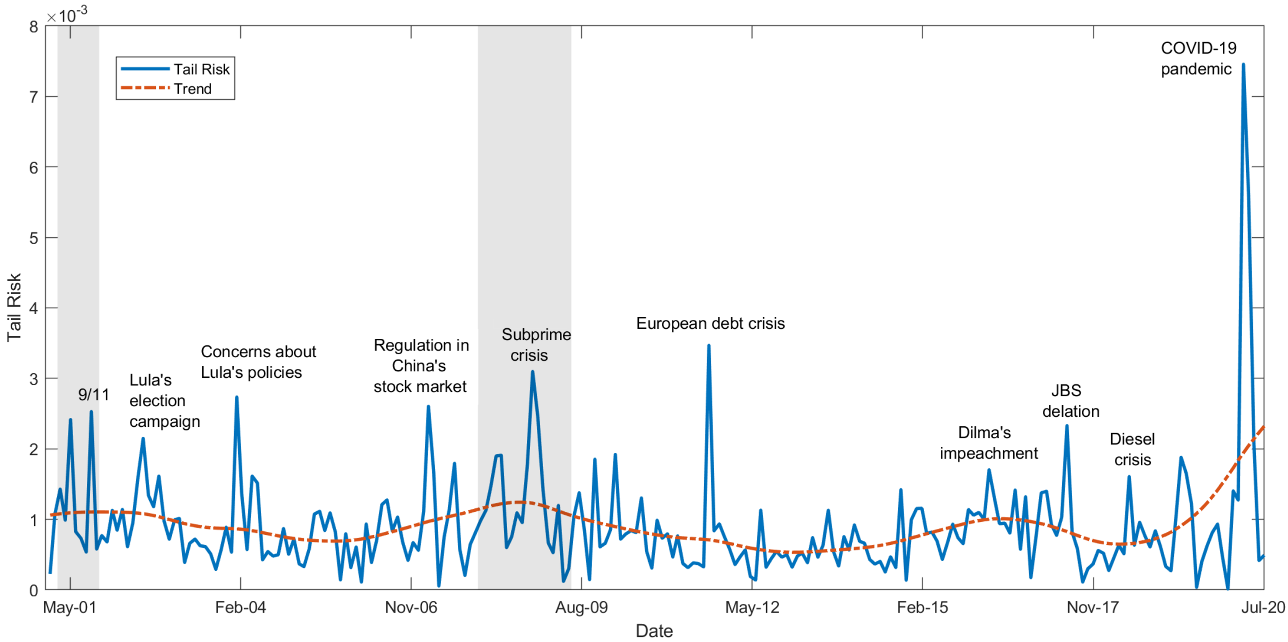

This figure plots the monthly tail risk measure in blue and the Hoddrick-Prescott (HP) trend in red dashed line, along with annotations corresponding to the most relevant tail risk peaks. The sample ranges from January, 2001 to July, 2020. Shaded areas depict NBER recession dates.

Figure 4.2: **Tail Risk and EPU (2001-2020)**

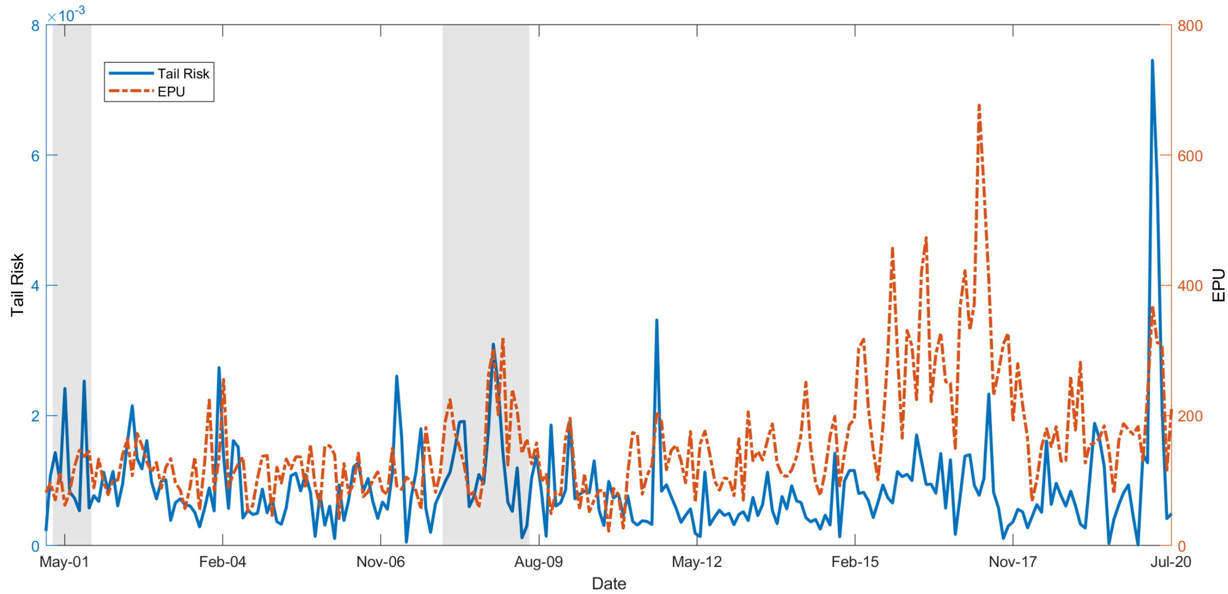

This figure plots the monthly tail risk measure in blue and the EPU for Brazil in red dashed line. The correlation between tail risk and EPU is 25.31%. The sample ranges from January, 2001 to July, 2020. Shaded areas depict NBER recession dates.

Figure 4.3: **Tail Risk and IVol-BR (2011-2020)**

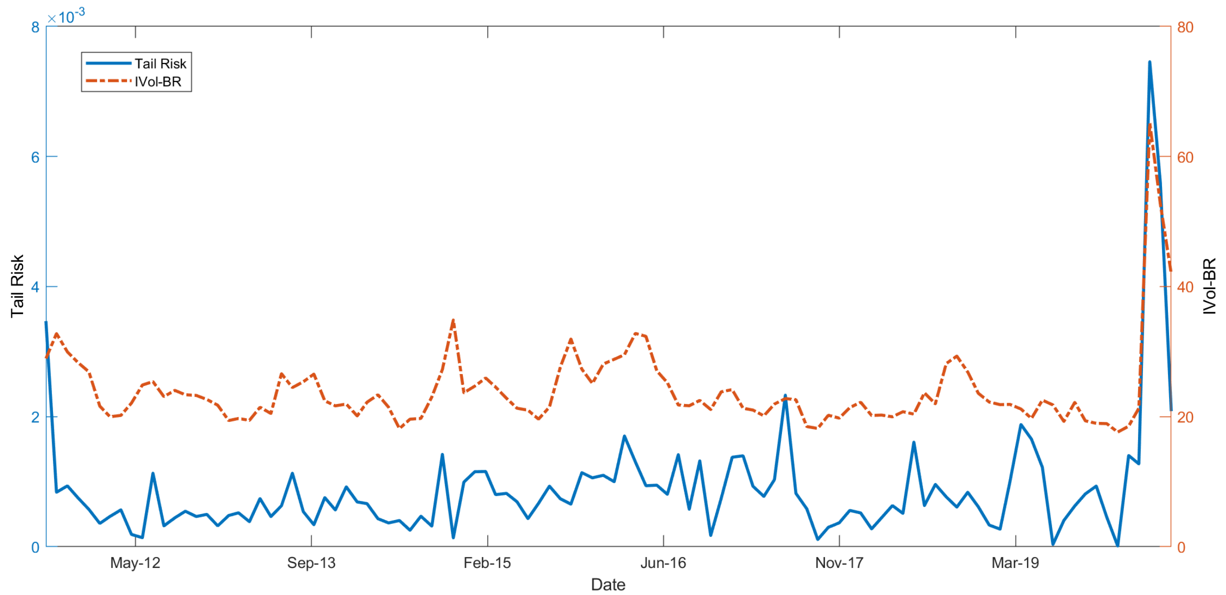

This figure plots the monthly tail risk measure in blue and the IVol-BR index from Astorino et al. (2017) in red dashed line. The correlation between tail risk and IVol-BR is 76.67%. The sample ranges from August, 2011 to July, 2020.

Figure 4.4: **Tail Risk Impulse Response Functions**

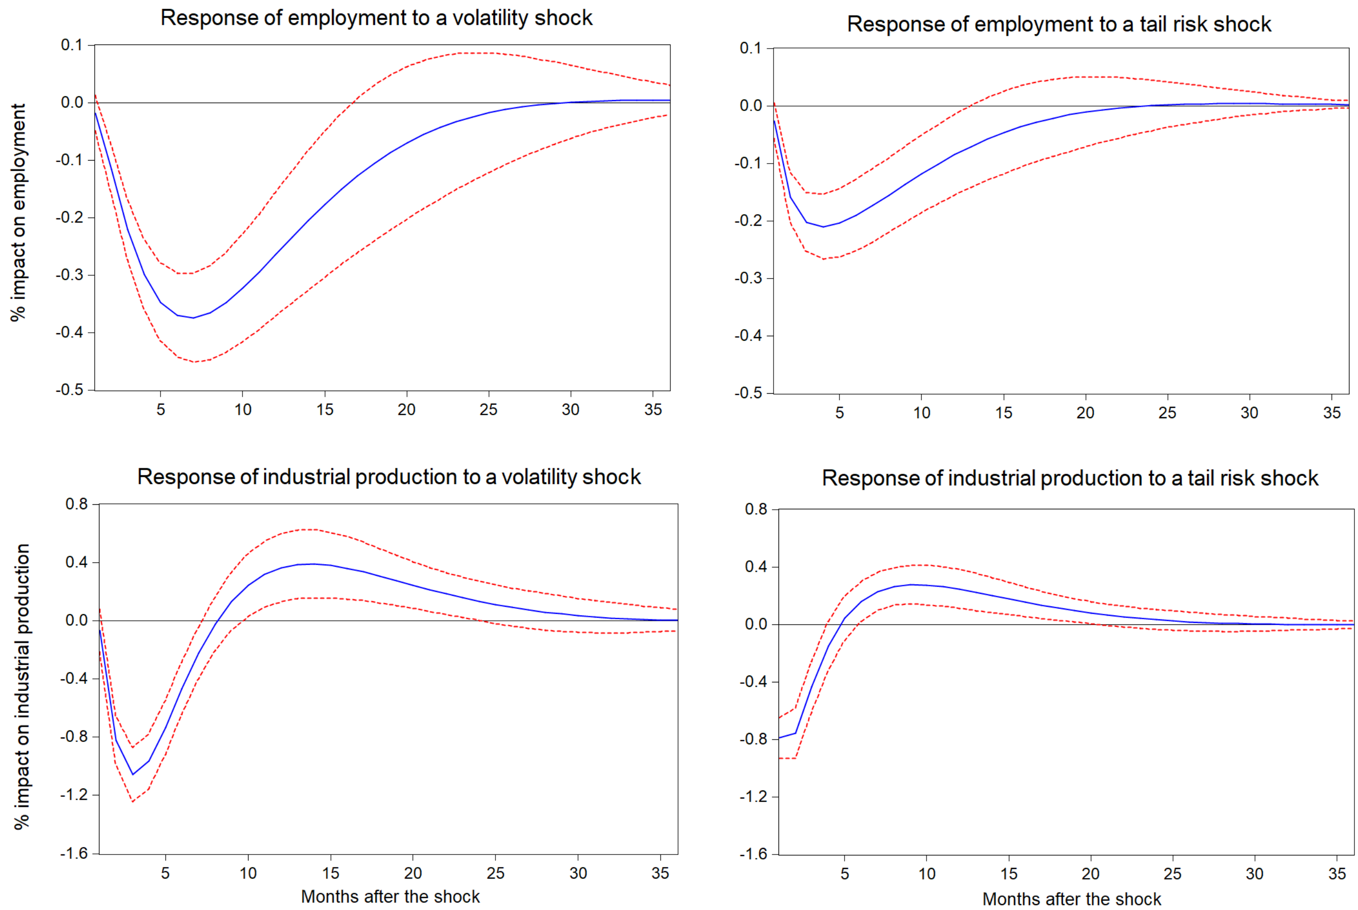

This figure plots the impulse responses for a one-standard-deviation shock to volatility and a one-standard-deviation shock to tail risk on employment and industrial production. The impact is estimated from a monthly VAR that includes log(Bovespa stock market index), stock market volatility, tail risk, log(employment) and log(industrial production) over the period March 2002 to July 2020. All variables are Hodrick-Prescott (HP) detrended (with $\lambda = 129,600$). Red dashed lines are one standard error bands following Bloom (2009). Vertical axis is in percent.

Figure 4.5: **Important Words for Tail Risk**

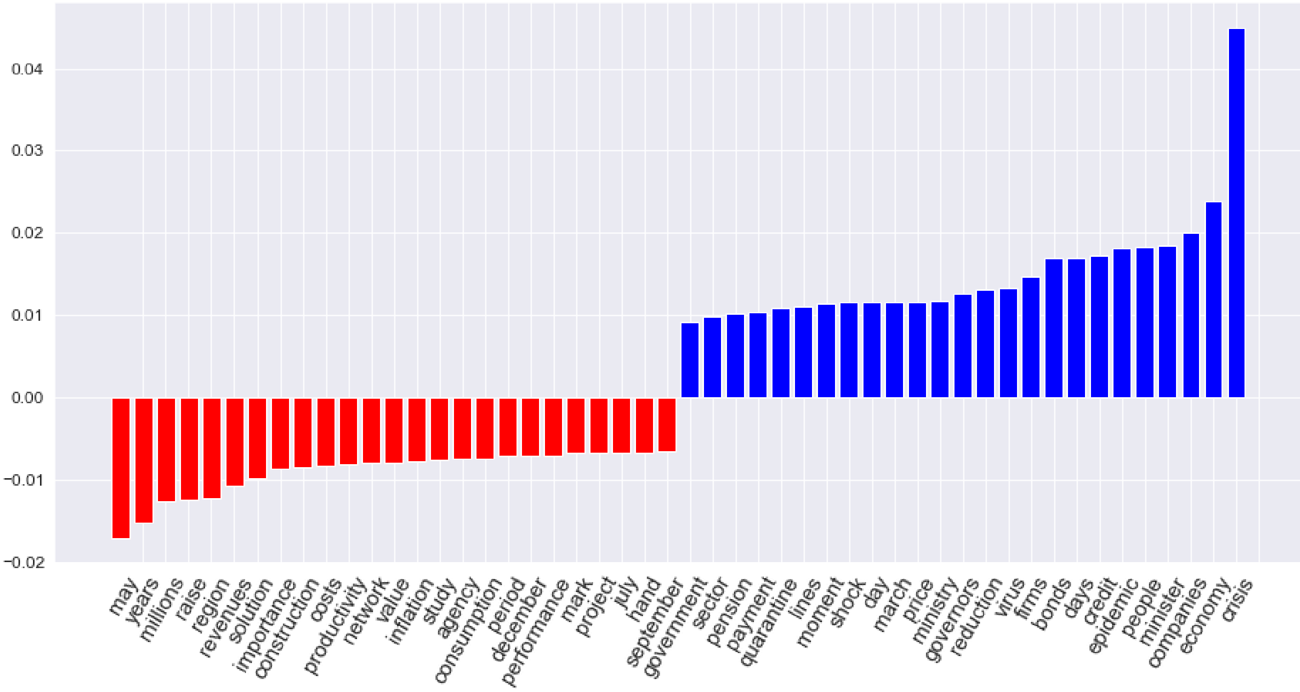

This figure plots the top 25 largest and smallest coefficients for tail risk and corresponding n-grams estimated by elastic net. The sample period is August, 2011 to July, 2020.

Figure 4.6: **Categories Implied Tail Risk**

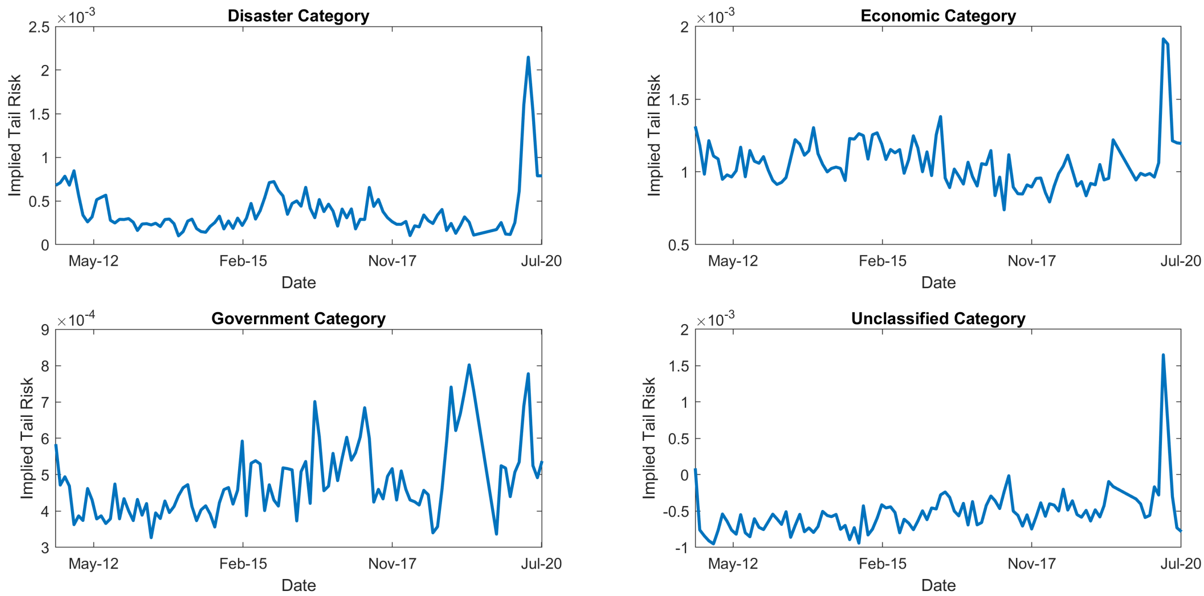

This figure plots the tail risk implied by the disaster, economic, government and unclassified categories. The sample ranges from August, 2011 to July, 2020.

Figure 4.7: **Categories Implied Tail Risk Impulse Response Functions**

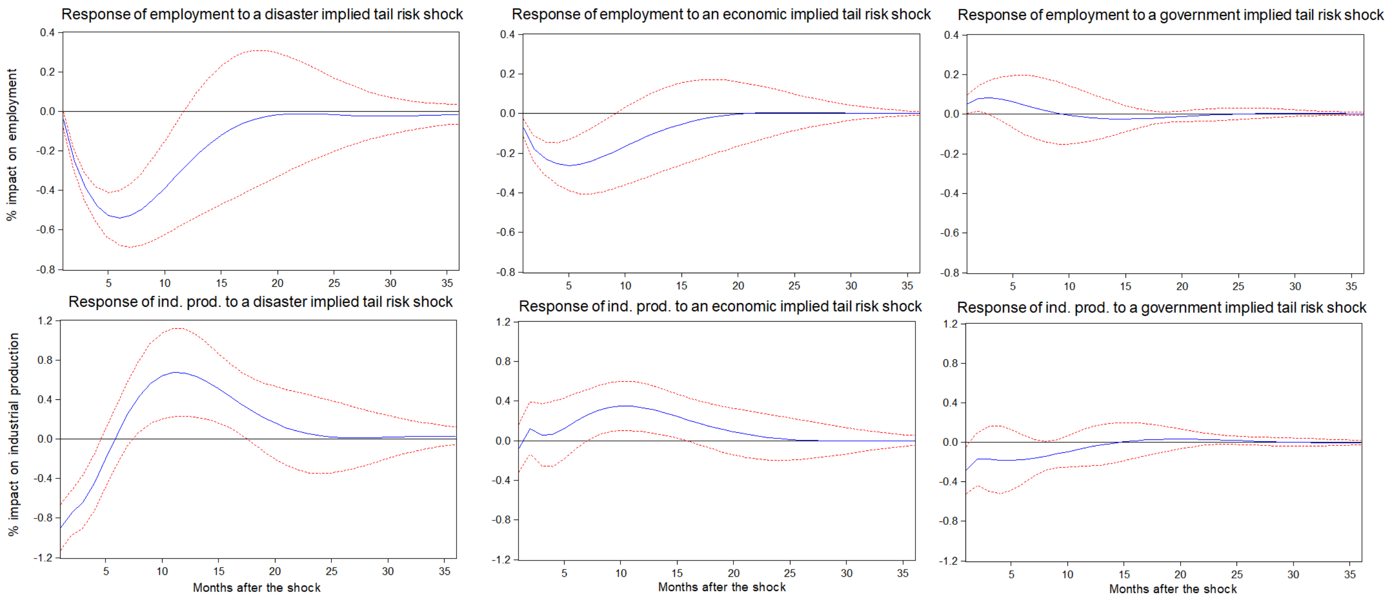

This figure plots the impulse responses for a one-standard-deviation shock to categories implied tail risk on employment and industrial production. The impact is separately estimated for each category, from a monthly VAR that includes log(Bovespa stock market index), stock market volatility, the category implied tail risk, log(employment) and log(industrial production) over the period August 2011 to July 2020. All variables are Hodrick-Prescott (HP) detrended (with $\lambda = 129,600$). Red dashed lines are one standard error bands following Bloom (2009). Vertical axis is in percent.

References

AIT-SAHALIA, Y.; LO, A. W. Nonparametric estimation of state-price densities implicit in financial asset prices. The Journal of Finance, v. 53, p. 499-547, 1998.

AIT-SAHALIA, Y.; LO, A. W. Nonparametric risk management and implied risk aversion. Journal of Econometrics, v. 94, p. 9-51, 2000.

AIT-SAHALIA, Y.; WANG, Y.; YARED, F. Do option markets correctly price the probabilities of movement of the underlying asset? Journal of Econometrics, v. 102, p. 67-110, 2001.

AIT-SAHALIA, Y.; DUARTE, J. Nonparametric option pricing under shape restrictions. Journal of Econometrics, v. 116, p. 9-47, 2003.

ALMEIDA, C.; ARDISON, K.; GARCIA, R.; VICENTE, J. Nonparametric tail risk, stock returns and the macroeconomy. Journal of Financial Econometrics, v. 15, p. 333-376, 2017.

ALMEIDA, C.; ARDISON, K.; GARCIA, R.; ORLOWSKI, P. Extracting tail risk from high-frequency S&P 500 returns. Unpublished working paper, 2020.

ALMEIDA, C.; GARCIA, R. Assessing misspecified asset pricing models with empirical likelihood estimators. Journal of Econometrics, v. 170, p. 519–537, 2012.

ALMEIDA, C.; GARCIA, R. Economic implications of nonlinear pricing kernels. Management Science, v. 63, p. 3147-3529, 2017.

ANTOINE, B.; BONNAL, H.; RENAULT, E. On the efficient use of the informational content of estimating equations: implied probabilities and Euclidean empirical likelihood. Journal of Econometrics, v. 138, p. 461–487, 2007.

ASTORINO, E. S.; CHAGUE, F.; GIOVANNETTI, B.; DA SILVA, M. E. Variance premium and implied volatility in a low-liquidity option market. Revista Brasileira de Economia, v. 71, p. 3-28, 2017.

BAGGERLY, K. Empirical likelihood as a goodness-of-fit measure. Biometrika, v. 85, p. 535–547, 1998.

BAKER, S. R.; BLOOM, N.; DAVIS, S. J. Measuring economic policy uncertainty. The Quarterly Journal of Economics, v. 131, p. 1593-1636, 2016.

BAKSHI, G.; CAO, C.; CHEN, Z. Empirical performance of alternative option pricing models. The Journal of Finance, v. 52, p. 2003-2049, 1997.

BAKSHI, G.; KAPADIA, N.; MADAN, D. Stock return characteristics, skew laws, and the differential pricing of individual equity options. The Review of Financial Studies, v. 16, p. 101–143, 2003.

BAKSHI, G.; MADAN, D. A theory of volatility spreads. Management Science, v. 52, p. 1945–1956, 2006.

BAKSHI, G.; MADAN, D. Investor heterogeneity and the non-monotonicity of the aggregate marginal rate of substitution in the market index. Unpublished working paper, University of Maryland, 2008.

BAKSHI, G.; MADAN, D.; PANAYOTOV, G. Returns of claims on the upside and the viability of u-shaped pricing kernels. Journal of Financial Economics, v. 97, p. 130–154, 2010.

BANSAL, R.; LEHMANN, B. N. Growth-optimal portfolio restrictions on asset pricing models. Macroeconomic Dynamics, v. 1, p. 333-354, 1997.

BARONE-ADESI, G.; ENGLE, R.; MANCINI, L. A GARCH option pricing model with filtered historical simulation. The Review of Financial Studies, v. 21, p. 1223–1258, 2008.

BARONE-ADESI, G.; FUSARI, N.; MIRA, A.; SALA, C. Option market trading activity and the estimation of the pricing kernel: a Bayesian approach. Journal of Econometrics, v. 216, p. 430-449, 2020.

BARRO, R. J. Rare disasters and asset markets in the twentieth century. The Quarterly Journal of Economics, v. 121, p. 823–866, 2006.

BATES, D. Post-'87 crash fears in the S&P 500 futures option market. Journal of Econometrics, v. 94, p. 181–238, 2000.

BEARE, B. K. Measure preserving derivatives and the pricing kernel puzzle. Journal of Mathematical Economics, v. 47, p. 689–697, 2011.

BEARE, B. K.; SCHMIDT, L. An empirical test of pricing kernel monotonicity. Journal of Applied Econometrics, v. 31, p. 338–356, 2014.

BERNARDO, A.; LEDOIT, O. Gain, loss and asset pricing. Journal of Political Economy, v. 108, p. 144–172, 2000.

BICKEL, P.; KLASSEN, C.; RITOV, Y.; WELLNER, J. Efficient and adaptive estimation for semiparametric models. Baltimore: Johns Hopkins Press, 1993.

BLACK, F.; SCHOLES, M. The pricing of options and corporate liabilities. Journal of Political Economy, v. 81, p. 637-654, 1973.

BLOOM, N. The impact of uncertainty shocks. Econometrica, v. 77, p. 623–85, 2009.

BOLLERSLEV, T.; TODOROV, V. Tails, fears, and risk premia. The Journal of Finance, v. 66, p. 2165-2211, 2011.

BOLLERSLEV, T.; TODOROV, V.; XU, L. Tail risk premia and return predictability. Journal of Financial Economics, v. 118, p. 113-134, 2015.

BORWEIN, J. M.; LEWIS, A. S. Duality relationships for entropy-like minimization problems. SIAM Journal on Control and Optimization, v. 29, p. 325–338, 1991.

BREEDEN, D.; LITZENBERGER, R. H. Prices of state-contingent claims implicit in option prices. Journal of Business, v. 51, p. 621-651, 1978.

BROADIE, M.; KAYA, O. Exact simulation of stochastic volatility and other affine jump diffusion processes. Operation Research, v. 54, p. 217–231, 2006.

BROADIE, M.; CHERNOV, M.; JOHANNES, M. Model specification and risk premia: evidence from futures options. The Journal of Finance, v. 62, p. 1453–1490, 2007.

BROGAARD, J.; DETZEL, A. The asset-pricing implications of government economic policy uncertainty. Management Science, v. 61, p. 3-18, 2015.

BROWN, B. W.; NEWEY, W. K. Efficient semiparametric estimation of expectations. Econometrica, v. 66, p. 453–464, 1998.

BROWN, B. W.; NEWEY, W. K. Generalized method of moments, efficient bootstrapping, and improved inference. Journal of Business and Economic Statistics, v. 20, p. 507–517, 2002.

BROWN, D.; JACKWERTH, J. C. The pricing kernel puzzle: reconciling index option data and economic theory. In: Contemporary Studies in Economics and Financial Analysis: Derivative Securities Pricing and Modelling, v. 94, p. 155–183. Bingley: Emerald Group, 2012.

CERNY, A. Generalized Sharpe ratios and asset pricing in incomplete markets. Review of Finance, v. 7, p. 191–233, 2003.

CHABI-YO, F.; GARCIA, R.; RENAULT, E. State dependence can explain the risk aversion puzzle. The Review of Financial Studies, v. 21, p. 973–1011, 2008.

CHEN, L.; PELGER, M.; ZHU, J. Deep learning in asset pricing. Unpublished working paper, 2019.

CHRISTOFFERSEN, P.; HESTON, S. L.; JACOBS, K. Capturing option anomalies with a variance-dependent pricing kernel. The Review of Financial Studies, v. 26, p. 1963–2006, 2013.

CHRISTOFFERSEN, P.; JACOBS, K.; ORNTHANALAI, C. GARCH Option Valuation: Theory and Evidence. The Journal of Derivatives, v. 21, p. 8-41, 2013.

CLARK, T. E.; MCCRACKEN, M. W. Tests of equal forecast accuracy and encompassing for nested models. Journal of Econometrics, v. 105, p. 85-110, 2001.

COCHRANE, J.; SAA-REQUEJO, J. Beyond arbitrage: good-deal asset price bounds in incomplete markets. Journal of Political Economy, v. 108, p. 79–119, 2000.

COCHRANE, J. Asset pricing. Princeton: Princeton University Press, 2001.

CONSTANTINIDES, G. M.; PERRAKIS, S. Stochastic dominance bounds on derivatives prices in a multiperiod economy with proportional transaction costs. Journal of Economic Dynamics and Control, v. 26, p. 1323-1352, 2002.

CONSTANTINIDES, G. M.; JACKWERTH, J. C.; PERRAKIS, S. Mispricing of S&P 500 index options. The Review of Financial Studies, v. 22, p. 1247–1277, 2009.

CRESSIE, N.; READ, T. Multinomial goodness-of-fit tests. Journal of the Royal Statistical Society Series B (Methodological), v. 46, p. 440-464, 1984.

CUESDEANU, H. Empirical pricing kernels: a tale of two tails and volatility. Unpublished working paper, University of Konstanz, 2016.

CUESDEANU, H.; JACKWERTH, J. C. The pricing kernel puzzle: survey and outlook. Annals of Finance, v. 14, p. 289-329, 2018.

DERMAN, E.; KANI, I. Riding on a smile. Risk, v. 7, p. 32-39, 1994.

DUAN, J. C. The GARCH option pricing model. Mathematical Finance, v. 5, p. 13-32, 1995.

DUAN, J. C.; SIMONATO, J. G. Empirical martingale simulation for asset prices. Management Science, v. 44, p. 1218-1233, 1998.

DUFFIE, D.; PAN, J.; SINGLETON, K. Transform analysis and asset pricing for affine jump-diffusions. Econometrica, v. 68, p. 1343–137, 2000.

DUFFIE, D. Dynamic asset pricing theory. Princeton: Princeton University Press, 2001.

DUPIRE, B. Pricing with a smile. Risk, v. 7, p. 18-20, 1994.

ERAKER, B.; JOHANNES, M.; POLSON, N. The impact of jumps in volatility and returns. The Journal of Finance, v. 58, p. 1269–1300, 2003.

FAMA, E. F.; FRENCH, K. R. A five-factor asset pricing model. Journal of Financial Economics, v. 116, p. 1-22, 2015.

FONSECA, E. R.; ROSA, J. L. G. Mac-morpho revisited: towards robust part-of-speech tagging. In: Proceedings of the 9th Brazilian Symposium in Information and Human Language Technology, Fortaleza, Brazil, 2013.

FUJIWARA, T.; MIYAHARA, Y. The minimal entropy martingale measures for geometric Lévy processes. Finance and Stochastics, v. 7, p. 509–531, 2003.

GABAIX, X. Variable rare disasters: an exactly solved framework for ten puzzles in macro-finance. The Quarterly Journal of Economics, v. 127, p. 645–700, 2012.

GARCIA, R.; GENÇAY, R. Pricing and hedging derivative securities with neural networks and a homogeneity hint. Journal of Econometrics, v. 94, p. 93-115, 2000.

GÂRLEANU, N; PEDERSEN, L. H.; POTESHMAN, A. M. Demand-based option pricing. The Review of Financial Studies, v. 22, p. 4259–4299, 2009.

GENTZKOW, M.; SHAPIRO, J. M. Media bias and reputation. Journal of Political

Economy, v. 114, p. 280-316, 2006.

GHOSH, A.; JULLIARD, C.; TAYLOR, A. An information-theoretic asset pricing model. Unpublished working paper, 2020.

GLOSTEN, L. R.; JAGANNATHAN, R.; RUNKLE, D. E. On the relation between the expected value and the volatility of the nominal excess return on stocks. The Journal of Finance, v. 48, p. 1779-1801, 1993.

GOURIO, F. Disaster risk and business cycles. American Economic Review, v. 102, p. 2734–2766, 2012.

GRAY, N.; NEWMAN, S. Canonical valuation of options in the presence of stochastic volatility. Journal of Futures Markets, v. 25, p. 1-19, 2005.

GU, S.; KELLY, B.; XIU, D. Empirical asset pricing via machine learning. The Review of Financial Studies, v. 33, p. 2223–2273, 2020.

HALEY, M. R.; WALKER, T. B. Alternative tilts for nonparametric option pricing. Journal of Futures Markets, v. 30, p. 983-1006, 2010.

HANSEN, L. P.; HEATON, J.; YARON, A. Finite-sample properties of some alternative GMM estimators. Journal of Business and Economic Statistics, v. 14, p. 262–280, 1996.

HANSEN, L. P.; JAGANNATHAN, R. Implications for security market data for models of dynamic economies. Journal of Political Economy, v. 99, p. 225–262, 1991.

HENS, T.; REICHLIN, C. Three solutions to the pricing kernel puzzle. Review of Finance, v. 17, p. 1065–1098, 2013.

HESTON, S. L. A closed-form solution for options with stochastic volatility with applications to bond and currency options. The Review of Financial Studies, v. 6, p. 327–343, 1993.

HESTON, S. L.; NANDI, S. A closed-form GARCH option valuation model. The Review of Financial Studies, v. 13, p. 585–625, 2000.

HUTCHINSON, J. M.; LO, A. W.; POGGIO, T. A nonparametric approach to the pricing and hedging of derivative securities via learning networks. The Journal of Finance, v. 49, p. 851-889, 1994.

JACKWERTH, J. C.; RUBINSTEIN, M. Recovering probability distributions from option prices. The Journal of Finance, v. 51, p. 1611-1631, 1996.

JACKWERTH, J. C. Recovering risk aversion from option prices and realized returns. The Review of Financial Studies, v. 13, p. 433–451, 2000.

JACKWERTH, J. C. Option-implied risk-neutral distributions and risk aversion, In: Research Foundation of AIMR. Charlottesville, VA, 2004.

KELLY, B.; JIANG, H. Tail risk and asset prices. The Review of Financial Studies, v. 27, p. 2841-2871, 2014.

KIMBALL, M. S. Precautionary saving in the small and in the large. Econometrica, v. 58, p. 53-73, 1990.

KITAMURA, Y. GNP-optimal tests for moment restrictions. Unpublished working paper, Department of Economics, University of Minnesota, 1996.

KITAMURA, Y. Empirical likelihood methods in econometrics: theory and practice. Cowles Foundation Discussion Papers 1569, Yale University, 2006.

KITAMURA, Y.; OTSU, T.; EVDOKIMOV, K. Robustness, infinitesimal neighborhoods, and moment restrictions. Econometrica, v. 81, p. 1185-1201, 2013.

KITAMURA, Y.; STUTZER, M. An information theoretic alternative to generalized method of moments estimation. Econometrica, v. 65, p. 861–874, 1997.

KORSAYE, S. A.; QUAINI, A.; TROJANI, F. Smart SDFs. Unpublished working paper, 2020.

KOZAK, S.; NAGEL, S.; SANTOSH, S. Shrinking the cross-section. Journal of Financial Economics, v. 135, p. 271-292, 2020.

LEVY, H. Upper and lower bounds of put and call option value: stochastic dominance approach. The Journal of Finance, v. 40, p. 1197-1217, 1985.

LINN, M.; SHIVE, S.; SHUMWAY, T. Pricing kernel monotonicity and conditional information. The Review of Financial Studies, v. 31, p. 493–531, 2018.

LIU, Y. Index option returns and generalized entropy bounds. Journal of Financial Economics, in press, 2020.

LOUGHRAN, T.; MCDONALD, B. When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. The Journal of Finance, v. 66, p. 35-65, 2011.

LUCAS, R. Asset prices in an exchange economy. Econometrica, v. 46, p. 1429–1445, 1978.

MANELA, A. Spreading information and media coverage: theory and evidence from drug approvals. Ph.D. thesis, University of Chicago, Chicago, 2011.

MANELA, A.; MOREIRA, A. News implied volatility and disaster concerns. Journal of Financial Economics, v. 123, p. 137-162, 2017.

MARTIN, I. What is the expected return on the market? The Quarterly Journal of Economics, v. 132, p. 367-433, 2017.

MENEZES, C.; GEISS, C.; TRESSLER, J. Increasing downside risk. American Economic Review, v. 70, p. 921-932, 1980.

MERTON, R. C. Theory of rational option pricing. Bell Journal of Economics and Management Science, v. 4, p. 141–183, 1973.

NEWEY, W. K.; SMITH, R. Higher order properties of GMM and generalized empirical likelihood estimators. Econometrica, v. 72, p. 219–255, 2004.

NEWEY, W. K.; WEST, K. D. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica, v. 55, p. 703-708, 1987.

OWEN, A. Empirical likelihood ratio confidence intervals for a single functional. Biometrika, v. 75, p. 237–249, 1988.

PERRAKIS, S.; RYAN, P. Option bounds in discrete time. The Journal of Finance, v. 39, p. 519-525, 1984.

RIETZ, T. A. The equity risk premium a solution. Journal of Monetary Economics, v. 22, p. 117–131, 1988.

RITCHKEN, P. H. On option pricing bounds. The Journal of Finance, v. 40, p. 1219–1233, 1985.

ROSENBERG, J.; ENGLE, R. Empirical pricing kernels. Journal of Financial Economics, v. 64, p. 341-372, 2002.

RUBINSTEIN, M. The valuation of uncertain income streams and the pricing of options. The Bell Journal of Economics, v. 7, p. 407-425, 1976.

RUBINSTEIN, M. Implied binomial trees. The Journal of Finance, v. 49, p. 771–818, 1994.

SCHNEIDER, P.; TROJANI, F. Fear trading. Swiss Finance Institute Research Paper, 15-03, 2015.

SMITH, R. J. Alternative semi-parametric likelihood approaches to generalized method of moments estimation. Economic Journal, v. 107, p. 503–519, 1997.

SMITH, R. J. GEL criteria for moment condition models. Unpublished working paper, University of Warwick, 2004.

SNOW, K. Diagnosing asset pricing models using the distribution of assets returns. The Journal of Finance, v. 46, p. 955–883, 1991.

SONG, Z.; XIU, D. A tale of two option markets: state-price densities and volatility risk. Journal of Econometrics, v. 190, p. 176–196, 2016.

STUTZER, M. A Bayesian approach to diagnosis of asset pricing models. Journal of Econometrics, v. 68, p. 367–397, 1995.

STUTZER, M. A simple nonparametric approach to derivative security valuation. The Journal of Finance, v. 51, p. 1633–1652, 1996.

TETLOCK, P. C. Giving content to investor sentiment: the role of media in the stock market. The Journal of Finance, v. 62, p. 1139-1168, 2007.

WACHTER, J. A. Can time-varying risk of rare disasters explain aggregate stock market volatility? The Journal of Finance, v. 68, p. 987–1035, 2013.

WELLER, B. M. Measuring tail risks at high frequency. The Review of Financial Studies, v. 32, p. 3571–3616, 2019.

ZHU, K.; LING, S. Model-based pricing for financial derivatives. Journal of Econometrics, v. 187, p. 447-457, 2015.

ZIEGLER, A. Why does implied risk aversion smile? The Review of Financial Studies, v. 20, p. 859–904, 2007.

ZOU, H.; HASTIE, T. Regularization and variable selection via the elastic net. Journal of the Royal Statistics Society Series B, v. 67, p. 301-320, 2005.