

WORKING PAPER · NO. 2021-71

Option-Implied Spreads and Option Risk Premia

Christopher L. Culp, Mihir Gandhi, Yoshio Nozawa, and Pietro Veronesi

JUNE 2021

OPTION-IMPLIED SPREADS AND OPTION RISK PREMIA

Christopher L. Culp
Mihir Gandhi
Yoshio Nozawa
Pietro Veronesi

June 2021

For their comments, we thank Gerardo Manzo, Dmitriy Muravyev, and seminar participants at Bocconi University, Leeds School of Business, University of Chicago, and the 2019 AFA meetings. Veronesi acknowledges financial support from the Fama-Miller Center for Research in Finance and the Center for Research in Security Prices at the University of Chicago Booth School of Business. The views expressed here are those of the authors and do not necessarily reflect the views of any institution with which the authors are affiliated. Culp passed away in June 2020. We dedicate this paper to his memory. He was a wonderful coauthor, colleague, and friend.

© 2021 by Christopher L. Culp, Mihir Gandhi, Yoshio Nozawa, and Pietro Veronesi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Option-Implied Spreads and Option Risk Premia
Christopher L. Culp, Mihir Gandhi, Yoshio Nozawa, and Pietro Veronesi
June 2021
JEL No. G12,G13

ABSTRACT

We propose implied spreads (IS) and normalized implied spreads (NIS) as simple measures to characterize option prices. IS is the credit spread of an option's implied bond, the portfolio long a risk-free bond and short a put option. NIS normalizes IS by the risk-neutral default probability and reflects tail risk. IS and NIS are countercyclical and predict implied bond returns, while neither, like implied volatility, predicts put returns. These opposite predictability results are consistent with a stochastic volatility, stochastic jump intensity model, as put premia increase in volatility but decrease in jump intensity, while implied bond premia increase in both.

Christopher L. Culp
Johns Hopkins Institute
for Applied Economics
Global Health and Study
of Business Enterprise
3400 North Charles St.
Baltimore, MD 21218
cculp@rmcsinc.com

Yoshio Nozawa
Lee Shau Kee Business Building
HKUST Business School
Clearwater Bay Road
Hong Kong, China
Sai Kung, NT
Hong Kong
ynchoir@gmail.com

Mihir Gandhi
University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
mihir.a.gandhi@gmail.com

Pietro Veronesi
University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
pietro.veronesi@chicagobooth.edu

1. Introduction

We introduce two novel quantities that characterize the relative pricing of financial options and that allow us to uncover the elusive nature of the dynamics of option risk premia. The first is an option’s implied spread (IS): this quantity builds on the seminal insight of Merton (1974) that a zero-coupon corporate bond is economically equivalent to safe debt minus a put option. We turn this insight on its head and from each traded put option we construct an option’s implied bond, a defaultable zero-coupon bond. The implied spread is the credit spread of the implied bond. The second quantity is the normalized implied spread, which normalizes the implied spread by its option-implied expected default frequency. We show that looking at options from a defaultable bond perspective, i.e. by exploiting the notion of implied bonds and implied spreads, allows us to uncover new dynamics of option risk premia and explain several regularities observed in the data.


To illustrate, consider the portfolio long a risk-free zero-coupon bond with unit face value and short $1/K$ European put options with strike price K and value $P(K, T)$. The value of this portfolio is $B(K, T) = Z(T) - P(K, T)/K$, where $Z(T)$ is the risk-free discount factor for horizon T . This portfolio pays \$1 if the underlying stock price is above K at maturity or the ratio of the stock price to the strike price otherwise. Because payoff structure of this portfolio is bond-like (as in Merton 1974), we refer to it as the option’s *implied bond*. Like any zero-coupon bond, the implied bond has an annualized yield $y(K, T) = -\log(B(K, T))/T$. The *implied spread* is the spread of the implied bond over risk-free Treasuries: $IS(K, T) = y(K, T) - r(T)$, where $r(T)$ is the risk-free rate for maturity T .

As an example, Figure 1 plots option-implied surfaces for S&P 500 index (SPX) options on October 6, 2017. Panel A plots implied spreads across the moneyness K/S and maturity T plane. Implied spreads are well-behaved: they increase with moneyness as market leverage and so default probabilities increase. These spreads measure the economic return on the investment in the bond and are comparable to standard corporate bond spreads (see Culp, Nozawa, and Veronesi 2018). For comparison, Panel C plots the underlying put option prices, and Panel D plots the corresponding Black-Scholes-Merton implied volatilities (IVs). As is well known, implied volatilities are not constant in the moneyness/maturity plane and instead exhibit the usual smirk across moneyness.

One shortcoming of the implied spread is that it increases to infinity for short-maturity, in-the-money (ITM) options as the probability the option will be in-the-money, or the implied bond will default, at maturity increases to one. To address this, we introduce our second measure for the relative pricing of options: we normalize the implied spread $IS(K, T)$ by the

Figure 1: Option-Implied Surfaces on October 6, 2017

This figure plots implied spreads (Panel A), normalized implied spreads (Panel B), put option prices (Panel C), and implied volatilities (Panel D) by moneyness and time-to-maturity on October 6, 2017.



implied bond's annualized, risk-adjusted default intensity $IDF(K, T)$. We refer to this as the option's *normalized implied spread* (NIS). We can compute default intensities immediately from neighboring put option prices, and so the normalization is in fact straightforward.

Panel B of Figure 1 plots the normalized implied spreads that correspond to the implied spreads in Panel A. The normalization indeed stabilizes spreads across strike prices and maturities. The normalized implied spread in Panel B and implied volatility surface in Panel D have somewhat similar shapes. Moreover, the misalignment in put option prices (violation of no-arbitrage) is easily visible from the bulge in normalized implied spreads among high strike prices. The misalignment is also evident from the put option prices in Panel C.

In this paper, we study the empirical properties of option-implied spreads and discuss their relation to the variation of option risk premia. Consistent with the corporate bond literature, we find that implied spreads strongly predict future implied bond returns. On the other hand, neither implied volatility (as is well known) nor implied spreads predict future option returns. We investigate these empirical regularities, discuss the relation between implied bond returns and option returns, and finally propose a two-factor option pricing model that reconciles our empirical findings.

More specifically, like IV, IS and NIS have appealing empirical properties. First, IS and NIS are countercyclical, increasing in recessions and decreasing in booms. This empirical regularity is similar to empirical findings in Culp, Nozawa, and Veronesi (2018), but we provide further evidence from the cross-section of option prices. Tail risk, in particular, is more acute during booms than during recessions.

Second, IS and NIS both predict excess returns on implied bonds, portfolios long safe bonds and short put options. This predictability highlights an important risk premium in implied spreads analogous to that in corporate bond spreads. In good times, implied spreads and risk premia are small, indicating low future returns on implied bonds. In bad times, implied spreads and risk premia are large, indicating high future returns on implied bonds. Our empirical results provide further evidence of a substantial time-varying risk premium as a compensation for tail risk, which characterizes the payoff structure of implied bonds.

Consistent with a time-varying premium for tail risk, we find that excess returns on out-of-the-money (OTM) implied bonds, which have lower default probabilities, are more predictable (larger t-statistics and higher R^2 s) than those on at-the-money (ATM) implied bonds, as tail risk affects the former more than the latter. This lines up with similar evidence for corporate bonds: the excess returns on investment-grade bonds are more predictable than those on high-yield bonds (see, for example, Nozawa 2017).

That implied spreads predict implied bond returns is consistent with a standard option pricing model with stochastic volatility and jumps (SVJ). In this model, higher IS and NIS imply higher implied bond risk premia. The model also indicates predictability is stronger for OTM implied bonds, as in the data.

IV, IS, and NIS, however, *do not* predict put option excess returns. This empirical finding is puzzling in light of the predictability of implied bond excess returns, as implied bonds are just risk-free bonds minus put options. Indeed, the standard SVJ model mentioned above also counterfactually implies that the put option risk premium is nonlinearly but positively related to IV, IS, and NIS. The reason is that in the SVJ model the only source of time-varying risk premium is time-varying volatility. But we find little evidence of a positive slope and only mild evidence of nonlinearity in the data. In sum, while the standard SVJ model is consistent with the evidence from implied bonds, it is at odds with the evidence from options themselves.

The counterfactual predictions of the one-factor SVJ model call for a two-factor model. We show that a two-factor, stochastic volatility, stochastic jump intensity (SVSIJ) model indeed rationalizes both implied bond and put option return predictability.

In the SVSIJ model, implied spreads, normalized implied spreads, and implied bond risk premia all increase with both volatility and intensity, and thus, IS and NIS naturally predict implied bond returns. However, volatility and intensity have opposite effects on put option risk premia: while risk premia increase with volatility (as in SVJ), risk premia decrease with intensity. Intuitively, put options are hedges against stock market crashes: as the probability of a jump increases, the hedge becomes more valuable, and so the risk premium decreases. Because of the asymmetric effect of volatility and intensity on put option risk premia, IV, IS and NIS do not predict put option returns. We simulate the SVSIJ model under our parameter estimates and confirm in artificial data that implied bond returns are predictable while put option returns are not, as in the empirical data.

Implied spreads and normalized implied spread also offer simple alternatives to implied volatilities to gauge the relative pricing of options. While implied volatility surfaces are widely used by both practitioners and researchers, they suffer from numerous shortcomings. First, IV depends on a specific model – the Black-Scholes-Merton option pricing model – and its computation requires the numerical solution of an equation. Second, IVs are not economically coherent across strike prices and maturities because the Black-Scholes-Merton model assumes that volatility is constant. Third, while implied volatility roughly captures the relative pricing of options across strike prices and maturities, it is difficult to quantify and interpret the economic magnitude. For instance, if one option trades at 20% volatility and another at 30%, the economic difference in these prices is unclear. Implied spreads, in contrast, are economically meaningful as they represent credit spreads of the implied bonds. Normalized implied spreads divide by the default frequency and thus have the interpretation of a return-to-risk ratio.

Finally, an additional contribution of the paper is to introduce a novel empirical methodology to estimate structural models, such as the SVSIJ model above. We exploit a maximum likelihood methodology paired with a set of moment conditions that are designed to emphasize the predictability of implied bond returns and the (lack of) predictability of put option returns. This technique allows us to obtain parameter estimates that not only fit the dynamics of options data, but also provide the proper conditional moments (return predictability) observed in the data. The estimation reveals a modest market price of default-intensity risk, but a large negative estimate of the risk-neutral mean jump. These parameter values induce the negative relation between jump intensity and put option risk premia, which countervail the positive relation between volatility and put option risk premia.

In sum, looking at options from the lens of their defaultable bonds counterpart, that is, as implied bonds, allows us to uncover new empirical regularities and better understand the

nature of the variation in put option risk premia, which has evaded the empirical literature so far.

Related Literature. Our work relates to several strands of literature. First, our paper builds on Culp, Nozawa, and Veronesi (2018, CNV henceforth), who study the empirical properties of option-implied spreads among deep OTM put options. The focus and the empirical strategy of CNV differs from our own. CNV examine the sizes and time-series properties of implied spreads – called pseudo spreads in their paper – and compare them to corporate bond spreads. The goal of their paper is to learn about the sources of credit spreads and the factors that drive their dynamics. CNV provide evidence that idiosyncratic tail risk is an important component of pseudo spreads. They also show that an index of pseudo spreads strongly predicts future economic growth. Given their focus is to compare pseudo bonds with corporate bonds, CNV only use long-maturity, deep OTM options. In contrast, in this paper, we consider the entire spectrum of options in the moneyness/maturity plane, we introduce normalized implied spreads, and we focus on their empirical characteristics, especially in relation to uncovering the elusive dynamics of option risk premia.

Second, our work relates to the existing literature on covered call returns. The payoff on a covered call – a position long the underlying stock and short an ATM call option – is equivalent to that on an implied bond – a portfolio long a safe bond and short a put option – by put-call parity. Indeed, our empirical evidence provides an explanation for the covered call anomaly, the empirical observation that the covered call generates high excess returns (see Israelov and Nielsen 2015 for a discussion). Because a covered call is equivalent to an implied bond, our empirical results show that these excess returns stem in part from a tail risk premium, as it does for corporate bonds.

Third, our paper relates to the literature that studies option return predictability and forecasting with option-implied information (see Christoffersen, Jacobs, and Chang 2013 for a survey). For instance, Israelov and Kelly (2017) use an elaborate Monte Carlo simulation to forecast the distribution of option returns, and Cao, Goyal, Xiao, and Zhan (2019) use singlename implied volatilities to forecast the cross-section of corporate bond returns. Unlike this literature, we use option-implied spreads as predictors and identify option risk premia from the difference between implied bonds and put options.

Finally, our paper builds on existing stochastic volatility and stochastic jump intensity models, such as Merton (1976), Heston (1993), Bates (2000, 2006, 2012), Duffie, Pan, and Singleton (2000), Pan (2002), Eraker (2004), Broadie, Chernov, and Johannes (2009), Santa-Clara and Yan (2010), Christoffersen, Jacobs, and Ornathanalai (2012), and David and

Veronesi (2014), among others. We do not attempt a survey of this extensive literature here. Although some of these papers document risk premia due to stochastic volatilities, jumps, and stochastic jump intensities, none of these papers study implied bonds or implied spreads.

The remainder of this paper proceeds as follows. Section 2 defines (normalized) implied spreads and interprets their properties in benchmark option pricing models. Section 3 documents the cross-sectional, time-series, and business-cycle properties of option-implied surfaces. Section 4 evaluates the predictive content of option-implied spreads for future implied bond and put option returns. Section 5 analyzes return predictability in dynamic option pricing models. Section 6 concludes. The appendix contains technical details and additional results omitted from the text.

2. Option-Implied Spreads

This section lays out the empirical methodology. We define implied spreads and normalized implied spreads, examine their behavior in quiet and turbulent periods, and interpret these dynamics within simple log-normal and jump diffusion environments.

2.1. Implied Spreads

The foundation of our analysis is the option-implied bond: the implicit defaultable zero-coupon bond given by the portfolio long a safe zero-coupon bond and short a European put option. The implied spread is the credit spread of the implied bond.

More specifically, consider a European put option with strike price K , maturity date T , and value $P_t(K, T)$. The portfolio long a safe zero-coupon bond with face value K and short a put option with strike price K has bond-like payoff at maturity T . The portfolio pays the strike price K if $S_T > K$ or the recovery value S_T if $S_T \leq K$. We divide by the face value K to normalize the payoff to unity and obtain the following:

Definition: For a European put option $P_t(K, T)$ with strike price K and maturity date T ,

(a) The **implied bond** of a put option is

$$B_t(K, T) = Z_t(T) - \frac{P_t(K, T)}{K} \quad (1)$$

where $Z_t(T)$ is a risk-free zero-coupon bond with maturity T .

(b) The **implied spread** of a put option is the credit spread of its implied bond

$$IS_t(K, T) = y_t(K, T) - r_t(T) \quad (2)$$

where $y_t(K, T) = -\log(B_t(K, T))/(T - t)$ is the continuously compounded yield of the risky zero-coupon bond and $r_t(T) = -\log(Z_t(T))/(T - t)$ is that of the risk-free zero-coupon bond.

Substituting the expression for $y_t(K, T)$ into (2), a straightforward algebraic manipulation yields a simple formula for the implied spread:

$$IS_t(K, T) = -\frac{1}{T - t} \log \left(1 - \frac{1}{Z_t(T)} \frac{P_t(K, T)}{K} \right) \quad (3)$$

For a put option, the implied spread is well-defined from weak no-arbitrage bounds on European put options: $P_t(K, T) < K Z_t(K)$. For a call option, we first use put-call parity to obtain the equivalent put option and then use (3) to compute its implied spread.

Like implied volatility, the implied spread in (3) is an annualized percent, but it has the clear interpretation of a bond's yield in excess of the risk-free rate. A high implied spread indicates that the put option is expensive. Any analysis we might conduct on corporate bonds and their credit spreads we can do with implied bonds and their implied spreads.

2.2. Normalized Implied Spreads

One drawback of the implied spread is that the spread increases to infinity for short-maturity, ITM options as the implied bond defaults at maturity almost surely. This behavior is apparent in Panel A of Figure 1.

To resolve this issue, we normalize the spread by the implied bond's probability of default. Rather than turn to corporate bond default frequencies, our methodology relies only on options for parsimony, and thus, we compute the implied default probability directly from the surface of put option prices. In particular, the risk-neutral probability that $S_T < K$ is:¹

$$\text{Prob}[S_T < K] = \frac{1}{Z_t(T)} \frac{dP_t(K, T)}{dK}$$

The annualized implied default frequency of the bond is

$$IDF_t(K, T) = -\frac{1}{T - t} \log(1 - \text{Prob}[S_T < K]) \quad (4)$$

¹As is well known, $P(K, T) = \int_0^K Z_t(T) \max(K - S, 0) f^*(S) dS$, where $f^*(S)$ is the risk-neutral density. Thus, $dP/dK = Z_t(T) \int_0^K f^*(S) dS = Z_t(T) \text{Prob}[S_T < K]$.

Intuitively, $Q_t(T) = 1 - \text{Prob}[S_T < K] < 1$ is the survival probability between t and T and thus, $IDF_t(K, T)$ is the average default intensity if default were to occur at any time between t and T . That is, we simply translate the default probability into an annualized average intensity in the same way we transform a zero-coupon bond return into an annualized yield. Although a risky zero-coupon bond can only default at T , its payment in the event of non-default also only occurs at T (as it is zero-coupon bond). And so just as the yield $y_t(K, T)$ is an annualized average rate of return of a single payoff at T , the $IDF_t(K, T)$ is an annualized average intensity of a single default event at T . Yields and default intensities are analogous concepts.²

Substituting, we obtain the simple expression

$$IDF_t(K, T) = -\frac{1}{T-t} \log \left(1 - \frac{1}{Z_t(T)} \frac{dP_t(K, T)}{dK} \right) \quad (5)$$

There is an elegant symmetry between the expressions for the implied spread in (3) and for the implied default frequency in (5). The expressions are identical except that the former is in levels P/K and the latter is in first differences dP/dK .

The calculation of dP/dK is straightforward from neighboring put option prices with strike prices K_1 and K_2 and $K_1 < K < K_2$:

$$\frac{dP_t(K, T)}{dK} \approx \frac{P_t(K_2, T) - P_t(K_1, T)}{K_2 - K_1}$$

To ensure implied default intensities are well-defined, we do not directly compute the slope from the surface of raw put option prices. From (5), the implied default frequency is well-defined for any put option from standard no-arbitrage bounds on European put options: $0 < \frac{1}{Z_t(T)} \frac{dP_t(K, T)}{dK} < 1$. Thus, as it is important that the put price is a monotone function of the strike price, we first impose this shape restriction on the price surface and then compute the slope. We describe the procedure in the appendix.

Definition: The *normalized implied spread* of a put option is the implied spread $IS_t(K, T)$ in (2) divided by the annualized implied default frequency $IDF_t(K, T)$ in (4)

$$NIS_t(K, T) = \frac{IS_t(K, T)}{IDF_t(K, T)} \quad (6)$$

Panel B of Figure 1 shows that the normalization indeed stabilizes implied spreads across strike prices and maturities. The normalized implied spread thus offers an economic measure

²To further the intuition, consider a model with stochastic intensity of default λ_s . The risk-neutral survival probability is $\Pr[\text{No Default before } T \mid \text{No Default before } t] = E_t^* \left[e^{-\int_t^T \lambda_s ds} \right]$. The average intensity of default is thus $\Lambda_t(T) = -\log(\Pr[\text{No Default before } T \mid \text{No Default before } t]) / (T - t)$.

of the cross-sectional variation in put option prices. The economic quantity is the implied bond's spread scaled by its annualized (risk-neutral) default intensity. As a point of reference, the normalized implied spread is like a Sharpe ratio but in yield, not return, units. Just as the Sharpe ratio is an excess return scaled by volatility risk, the normalized implied spread is a credit spread scaled by default risk.³

An equivalent formulation of the normalized implied spread is

$$NIS_t(K, T) = \frac{\log(1 - LGD \times E_f^*[1_{x_T < 1}])}{\log(1 - E_f^*[1_{x_T < 1}])} \quad (7)$$

where $x_T = S_T/K$, $LGD = E_f^*[1 - x_T \mid x_T < 1]$, and $E_f^*[\cdot]$ denotes the expectation under the forward risk-adjusted probability measure. Thus, $E_f^*[1_{x_T < 1}]$ is the risk-adjusted default probability, LGD is the risk-adjusted expectation of loss-given-default, and so the numerator and denominator differ only in the size of LGD . It is clear from (7) that NIS is monotonically related to the risk-adjusted LGD . If $LGD = 0\%$, then $NIS = 0\%$. If $LGD = 100\%$, then $NIS = 100\%$. Indeed, we have the following proposition.

Proposition 1: *The normalized implied spread in (7) has the following properties:*

- (a) *NIS is increasing in the loss-given default LGD for a given $E_f^*[1_{x_T < 1}]$.*
- (b) *NIS is decreasing in the risk-adjusted default probability $E_f^*[1_{x_T < 1}]$ for a given LGD .*


Part (a) shows that NIS measures the size of an implied bond's (risk-adjusted) tail risk. For a given default probability, a high NIS implies the LGD is large.

Part (b) shows that an increase in the default probability *decreases* NIS . Mechanically, this follows directly from the definition. As $E_f^*[1_{x_T < 1}]$ converges to one, the denominator diverges to negative infinity while the numerator converges to a negative finite number. Thus, NIS converges to zero. Intuitively, the LGD – the size of tail risk – becomes *relatively* smaller the higher the default probability. If the default probability increases but LGD remains the same, then NIS declines. On the other hand, the default probability and LGD are positively correlated in practice, and thus, NIS may as well increase.

³To further the intuition, consider a simple model with constant default probability p , constant market price of risk π , and constant loss-given-default LGD . Then, the credit spread is $s = p\pi LGD$, and the risk-neutral default intensity is $p^* = p\pi$. It follows that $NIS = s/p^* = LGD$.

Figure 2: Option-Implied Quantities in Quiet and Turbulent Markets

This figure plots implied volatilities (top row), implied spreads (middle row), and normalized implied spreads (bottom row) by moneyness in the Financial Crisis (left column) and Covid-19 Crash (right column). Options have 60 days-to-maturity.



2.3. Case Study in Quiet and Turbulent Markets

To gauge the informational content of implied spreads and normalized implied spreads, we examine their dynamics in the Financial Crisis and the more recent Covid-19 Crash.

Financial Crisis. We begin with the 2008 Financial Crisis on the left side of Figure 2. Panel A plots implied volatilities in January 2007 and January 2009. The former is the calm before the storm and represents quiet markets: the average implied volatility hovers around 10%. The latter is the midst of the crisis and represents turbulent markets: the average implied volatility soars to 40%. That is, the failure of Lehman Brothers and the ensuing financial turmoil increased hedging demand for put options, and so implied volatility likewise

increased.

Panel B plots implied spreads. In 2007, implied spreads are low, between 1% for OTM options and 35% for ITM options. ITM options have high spreads because their implied bonds are extremely likely to default. Implied spreads too rise in the crisis: in 2009, OTM implied spreads jump to 25% and ITM implied spreads to 65%. Unlike implied volatility, the economic interpretation of a 25% spread is clear: it is the promised yield on a defaultable bond in excess of the risk-free rate.

Panel C shows the analogous impact of the crisis on normalized implied spreads. Like both implied volatilities and implied spreads, normalized implied spreads increase during the crisis. Because the normalized implied spread scales the implied spread by the (annualized) intensity of default, the increase between 2007 and 2009 suggests a substantial increase in tail risk in spite of the increase in default probability.

Covid-19 Crash. We next turn to the 2020 Covid-19 Crash on the right side of Figure 2. We compare option-implied quantities on February 19, the market peak before the crash, with those on March 20, the market trough before the recovery. Not surprisingly, implied volatilities, implied spreads, and normalized implied spreads largely rise, as in the Financial Crisis.⁴ The concavity of normalized implied spreads along strike prices also increases. As we see in Section 3, these properties not only characterize the Financial Crisis and Covid-19 Crash but also bad times more generally.

2.4. Benchmark Option Pricing Models


To better understand the dynamics of spreads in quiet and turbulent markets, we turn to two simple option pricing models.

Black-Scholes-Merton Model. The first benchmark is the canonical Black-Scholes-Merton model with constant volatility and no jumps in returns. Under the assumption stock return dynamics are log-normal $\log(S_T/S_0) \sim \mathcal{N}\left((r - \delta - \frac{1}{2}\sigma^2)T, \sigma^2T\right)$, the implied

⁴From peak to trough, implied spreads increase among OTM and ATM options, but decrease among ITM options. On February 19, volatility is low, the probability an ITM option is in-the-money at maturity is relatively high, and so the implied spread is relatively high. On March 20, volatility is high, the probability an ITM option is in-the-money at maturity is relatively low, and so the implied spread is relatively low.

Figure 3: Option-Implied Spreads in the Black-Scholes-Merton Model

This figure plots implied spreads (Panel A) and normalized implied spreads (Panel B) by moneyness in the Black-Scholes-Merton model for diffusive volatility $\sigma = 0.05, 0.15, 0.25$. Parameters: $r = 0.03$, $\delta = 0.02$, $T = 0.25$.



spread and normalized implied spread are

$$IS(M, T) = -\frac{1}{T} \log \left(1 - N(-d_2) + \frac{N(-d_1)}{M} \right)$$

$$NIS(M, T) = \frac{\log(1 - N(-d_2) + N(-d_1)/M)}{\log(1 - N(-d_2))}$$

where $M = K / (S_0 e^{(r-\delta)T})$ is the option's forward moneyness, $N(\cdot)$ is the cumulative normal density, and

$$d_1 = \frac{\log(S/K) + (r - \delta + \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}, \quad d_2 = d_1 - \sigma\sqrt{T}$$

are the usual constants.⁵


Panel A of Figure 3 plots implied spreads for three levels of volatility. Under Black-Scholes-Merton, implied spreads increase with moneyness and volatility. As moneyness (i.e. market leverage) and volatility increase, the probability of default also increases, and so implied spreads rise.

Panel B plots normalized implied spreads. In the model, normalized implied spreads

⁵The Black-Scholes-Merton dual delta is $dP/dK = Z(T)N(-d_2)$. With this, we compute normalized implied spreads in the usual way.

Figure 4: Option-Implied Spreads in Merton's Jump Diffusion Model

This figure plots implied spreads (top row) and normalized implied spreads (bottom row) by moneyness in Merton's jump diffusion model for diffusive volatility $\sigma = 0.10, 0.15$ and jump volatility $\sigma_J = 0.05, 0.15, 0.25$. Parameters: $r = 0.03$, $\delta = 0.02$, $T = 0.25$.



increase with and are concave in moneyness. But in the data, normalized implied spreads are in fact largely convex in moneyness, as seen in Figure 2 and as we show in Section 3. Intuitively, normalized implied spreads measure expected losses given default, but these losses are much too small under the assumption of log-normal dynamics. In other words, the Black-Scholes-Merton model does not capture the fat tails of the risk-neutral distribution (as is well known) and so cannot match the convexity of normalized implied spreads.

Merton's Jump Diffusion Model. The second benchmark is the Merton (1976) jump diffusion model. Under the assumption that log-normal jumps $\log(J) \sim \mathcal{N}(\mu_J, \sigma_J^2)$ arrive with (risk-neutral) Poisson intensity λ , the put option price in closed form is

$$Put^{MJ}(K, T) = \sum_{n=0}^{\infty} p_n \times Put^{BSM}(S_0, K, r_n, T, \delta, \sigma_n^2)$$

where $m = \mu_J + \frac{1}{2}\sigma_J^2$, $\lambda' = \lambda e^m$, $r_n = r - \lambda(e^m - 1) + \frac{n}{T}m$, and $\sigma_n^2 = \sigma^2 + \frac{n}{T}\sigma_J^2$. $p_n = \frac{1}{n!}e^{-\lambda'T}(\lambda'T)^n$ is the probability n jumps arrive over the life of the put option, and $Put^{BSM}(\cdot)$ is the corresponding Black-Scholes-Merton price. That is, with jumps, the put option price is the probability weighted sum of conditional prices. With put option prices, we compute (normalized) implied spreads in the usual way.

Figure 4 plots implied spreads and normalized implied spreads for two levels of diffusive volatility and three levels of jump volatility σ_J . While the impact of jumps on implied spreads is negligible, the impact on normalized implied spreads is substantial. This is unsurprising because jumps fatten the tails of the risk-neutral distribution, and normalized implied spreads measure tail risk. For OTM options, normalized implied spreads increase with σ_J , and the increase is larger the lower the diffusive volatility. Intuitively, the relative effect of the jump component is larger the lower diffusive volatility. For ITM options, normalized implied spreads also increase with σ_J , but the increase is small. Intuitively, the probability the implied bond defaults is so large that jumps have little impact on spreads (see Proposition 1). In short, the addition of jumps to the Black-Scholes-Merton model generates sufficient tail risk to qualitatively match the convexity of normalized implied spreads in the data.

3. Empirical Properties of Option-Implied Spreads

This section documents the cross-sectional and time-series properties of implied spreads and normalized implied spreads.


3.1. Data

Our focus is implied bonds on S&P 500 index put options. Option prices for the 1990 to 1995 period are from CBOE Market Data Express. Option prices for the 1996 to 2020 period are from OptionMetrics Ivy DB. Stock prices are from the Center for Research in Security Prices (CRSP). We use Treasury Bill rates and constant maturity Treasury Yields from the Federal Reserve Economic Database (FRED) to construct risk-free zero-coupon bonds. We convert these into zero-coupon yields and linearly interpolate to match option maturities. To filter clear violations of no-arbitrage and other quotation errors, we follow Constantinides, Jackwerth, and Savov (2013). We enumerate the exact filters in the appendix.

The main input for the empirical analysis is a panel of put option and implied bond portfolios with target moneyness $K/S = 0.90, 0.925, 0.95, 0.975, 1.00, 1.025, 1.05, 1.075, 1.10$ and target maturity $T = 30, 60, 91, 122, 152, 182, 273, 365$ days from January 1990 to December 2020. We discuss construction of this panel in the appendix.

Figure 5: Average Option-Implied Surfaces

This figure plots average implied volatilities (Panel A), average implied spreads (Panel B), average normalized implied spreads (Panel C), and average implied default intensities (Panel D) by moneyness and time-to-maturity. The sample is monthly from January 1990 to December 2020.



3.2. Cross-Sectional Properties


Unconditional Properties. Figure 5 plots unconditional average option-implied surfaces. Panel A is the standard implied volatility surface, which features the usual decreasing smirk across moneyness. Panel B shows the implied spread, which is instead increasing in moneyness (i.e. market leverage) as the default probability increases.

Panel C plots the normalized implied spread, the ratio of the implied spread in Panel B to the implied default intensity in Panel D. Like implied spreads, the normalized implied spread is mildly increasing across moneyness, although significantly less so than the implied spread. Unlike implied spreads, the normalized implied spread is increasing across maturities, which indicates that longer-term bonds have higher risk-adjusted credit spreads. Most important, the average normalized implied spread is convex with respect to moneyness, which reflects the tail risk embedded in put options (see discussion in Section 2.4).

Panel D shows the implied default frequency. Like implied spreads in Panel B, the implied default frequency is increasing in moneyness and decreasing in maturities. That is, the likelihood the bond defaults at maturity is highest for short-maturity, levered bonds. Equivalently, the likelihood the option ends in-the-money at maturity is highest for short-maturity,

Figure 6: Average Option-Implied Surfaces in Booms and Recessions

This figure plots average implied volatilities (Panel A), average implied spreads (Panel B), average normalized implied spreads (Panel C), and average implied default intensities (Panel D) by moneyness and time-to-maturity in booms (orange) and recessions (blue). The sample is monthly from January 1990 to December 2020.



ITM options. Neither interpretation is surprising. The similarity between the implied spread surface and implied default frequency surface reinforces the symmetry between their formulas in (3) and (5), respectively.

Business-Cycle Properties. Figure 6 plots average option-implied surfaces in booms and recessions. The difference between recessions and booms is largely positive along the full moneyness/maturity plane, as implied volatilities, implied spreads, normalized implied spreads, and implied default frequencies are higher in bad times. As is such, Panel A shows that implied volatilities increase across the board in recessions and especially so for short-maturity options. The same holds for implied spreads in Panel B. Short-maturity implied spreads increase the most in recessions.

Panel C shows the impact on normalized implied spreads. Like implied spreads, normalized implied spreads increase in recessions. Unlike implied spreads, the increase is greatest among long-maturity, OTM options. Intuitively, while implied default intensities (Panel D) increase in recessions, so do implied spreads (Panel B): at long maturities, the latter outweighs the former, and so normalized implied spreads increase.

In contrast, implied default intensities decrease in recessions among short-maturity, ITM options. In booms, volatility is low, the probability an ITM option is in-the-money at maturity is high, and so the implied default intensity is relatively high. In recessions, volatility is high, the probability an ITM option is in-the-money at maturity is low, and so the implied default intensity is relatively low. In other words, the same risk premium that drives implied spreads higher in recessions also drives implied default intensities and especially so at short maturities.

3.3. Time-Series Properties

Figure 7 plots the time-series of option-implied quantities. On the left are the time-series across strike prices for a single maturity $T = 2$ months. On the right are the time-series across maturities for a single moneyness $K/S = 1.00$. First, Panels A and B show the time-series of implied volatilities. As is well known, implied volatilities increase in bad times.

Second, Panels C and D show the time-series of implied spreads. Like implied volatilities, implied spreads increase in recessions and decrease in booms. The cross-sectional difference between spreads on ITM options (higher spreads) and OTM options (lower spreads) in Panel C is unconditionally large, which reflects the large default premium (see earlier discussion). In contrast, the cross-sectional difference across maturities in Panel D is relatively small.


Third, Panels E and F further highlight the time-series variation in normalized implied spreads. Since normalized implied spreads adjust for the default intensity, their time-series variation is akin to variation in the market price of risk implicit in put options. Unlike implied spreads, the cross-sectional variation in normalized implied spreads across strike prices is small, but across maturities it is large (see earlier discussion).

Finally, the cross-sectional variation over time is most interesting for implied volatilities across maturities (Panel B) and normalized implied spreads across strike prices (Panel E). The implied volatility term structure slopes upward in good times and downward in bad times, which reflects mean reversion. The dynamics of normalized implied spreads are less straightforward but in short reflect the time-varying role of tail risk (we dig more deeply into tail risk later). The cross-sectional variation for the remainder of the cuts is largely time invariant, and so their time variation is uninteresting.

To summarize the time variation in surface dynamics, Figure 8 plots the first principal component of each option-implied surface. The first principal component of implied volatilities and implied spreads nearly coincide, and just under this is that of normalized implied

Figure 7: Time-Series of Option-Implied Quantities

This figure plots the time-series of implied volatilities (top row), implied spreads (middle row), and normalized implied spreads (bottom row) by moneyness (left column) and time-to-maturity (right column). Grey bands are NBER recessions. The sample is monthly from January 1990 to December 2020.




spreads. The following table shows that the correlations among these principal components are large, both in levels and first differences. The lone exception is the correlation between implied spreads and normalized implied spreads in first differences, which is only 33% but still positive. These correlations further support the strong comovement among the time-series dynamics shown in Figure 7.

	Levels			First Differences	
	IS	NIS		IS	NIS
IV	0.92	0.84	IV	0.82	0.66
NIS	0.79		NIS	0.33	

Finally, Figure 9 plots the convexity of normalized implied spreads in blue. The proxy for

Figure 8: First Principal Component of Option-Implied Quantities

This figure plots the first principal component of implied volatilities (blue), implied spreads (red), and normalized implied spreads (yellow) in levels. Grey bands are NBER recessions. The sample is monthly from January 1990 to December 2020.



convexity is the butterfly spread across options with 30 days-to-maturity: $NIS(0.90, 30) - 2 \times NIS(1.00, 30) + NIS(1.10, 30)$. Over the sample, convexity is largely positive, in contrast with the negative prediction from the Black-Scholes-Merton model (see discussion in Section 2.4). Interestingly, normalized implied spreads are especially convex in good times but concave or near concave in bad times. For instance, the convexity is negative in every recession and in crises (e.g. the LTCM crisis in 1998 and the debt-ceiling crisis in 2013).


Figure 9 also plots the CBOE Skew index, the skewness of the risk-neutral distribution implied by SPX options, in orange.⁶ There is a clear association between convexity and skewness: the correlation between the two is 27% and significantly different from zero (not tabulated). The positive correlation furthers the intuition that normalized implied spreads capture some tail risk embedded in put option prices. In good times, returns are more negatively skewed, tail risk is high, and normalized implied spreads are more convex. In bad times, returns are less negatively skewed, tail risk is low, and normalized implied spreads are less convex.

In short, there are three key properties of option-implied spreads. First, (normalized) implied spreads are countercyclical, low in booms and high in recessions. Second, implied spreads increase with moneyness as market leverage and so the default probability increase. Finally, the convexity of normalized implied spreads is procyclical: convexity is high in booms when tail risk is high and low in recessions when tail risk is low.

⁶See <https://www.cboe.com/micro/skew/documents/skewwhitepaperjan2011.pdf> for more details.

Figure 9: Convexity of Normalized Implied Spreads

This figure plots the convexity of normalized implied spreads (blue) and the CBOE Skew index (orange). The convexity of normalized implied spreads is the butterfly spread $NIS(0.90, 30) - 2 \times NIS(1.00, 30) + NIS(1.10, 30)$, where $NIS(M, H)$ is the normalized implied spread of a put option with moneyness $M = K/S$ and H days-to-maturity. The convexity is -4.2% in March 1998 and -13.0% in February 2009. The CBOE Skew index is the 30-day ahead risk-neutral expectation of S&P 500 index skewness: a higher skew index indicates more negative skewness. The skew index is standardized to have mean zero and variance one. Grey bands are NBER recessions. The sample is monthly from January 1990 to December 2020.



4. Time-Varying Risk Premia

This section evaluates the predictive content of option-implied spreads for put option and implied bond excess returns. As is usual, the economic content of return predictability is evidence of time-varying risk premia. We begin with the one-month horizon and then turn to long horizons.

4.1. Monthly Returns

Our empirical analysis considers three excess returns, two for put options and one for implied bonds. The first is the put option excess return, the monthly return to a long put option position in excess of the risk-free rate:

$$R_{t+1}^{Put}(K, T) - R_t^f = \frac{P_{t+1}(K, T) - P_t(K, T)}{P_t(K, T)} - R_t^f \quad (8)$$

The second is the delta-hedged put option return, the monthly return to a long put option position in excess of the delta-hedged stock position:

$$R_{t+1}^{Put}(K, T) - \Delta_t R_{t+1}^{SPX} = \frac{P_{t+1}(K, T) - P_t(K, T)}{P_t(K, T)} - \Delta_t R_{t+1}^{SPX} \quad (9)$$

Table 1: Monthly Returns: Summary Statistics

This table reports summary statistics for put option excess returns $R^{Put} - R^f$, delta-hedged put option returns $R^{Put} - \Delta R^{SPX}$, and implied bond excess returns $R^{IB} - R^f$ by moneyness and time-to-maturity. OTM is moneyness $K/S \leq 0.95$, ATM is $0.975 \leq K/S \leq 1.025$, and ITM is $K/S \geq 1.05$. The sample is monthly from January 1990 to December 2020.

	$R^{Put} - R^f$		$R^{Put} - \Delta R^{SPX}$		$R^{IB} - R^f$	
	$T \leq 122$ days	$T \geq 152$ days	$T \leq 122$ days	$T \geq 152$ days	$T \leq 122$ days	$T \geq 152$ days
Panel A: Average Excess Return (Percent)						
OTM	-37.05	-10.75	-36.69	-10.32	0.42	0.43
ATM	-20.64	-7.97	-20.04	-7.42	0.64	0.52
ITM	-9.43	-5.61	-8.60	-4.88	0.70	0.62
Panel B: Standard Deviation (Percent)						
OTM	81.03	46.35	80.37	45.39	1.33	1.77
ATM	76.83	38.63	75.19	36.93	2.47	2.34
ITM	48.29	31.69	45.13	29.06	3.66	3.10
Panel C: Sharpe Ratio						
OTM	-0.46	-0.23	-0.46	-0.23	0.32	0.24
ATM	-0.27	-0.21	-0.27	-0.20	0.26	0.22
ITM	-0.20	-0.18	-0.19	-0.17	0.19	0.20

where Δ_t is the Black-Scholes-Merton delta. The delta-hedge ensures stock price dynamics do not drive option return dynamics. The third is the implied bond excess return, the monthly return to a long implied bond position in excess of the risk-free rate:

$$R_{t+1}^{IB}(K, T) - R_t^f = \frac{B_{t+1}(K, T) - B_t(K, T)}{B_t(K, T)} - R_t^f \quad (10)$$

Table 1 reports the mean, standard deviation, and Sharpe ratio for each of the returns. As is well known, put options earn negative returns on average, as they are levered, short positions in the stock. Returns are more negative among longer-maturity, deeper OTM options. The hedged put option return exhibits properties similar to the unhedged return. That is, at the monthly horizon, there is little impact of delta hedging on the return.

In contrast, the excess return on implied bonds is positive and an order of magnitude smaller than that on put options. Consistent with their bond-like payoff at maturity, implied bonds behave less like options and more like corporate bonds. OTM implied bonds, which have low default probability, earn smaller average returns than ITM implied bonds, which have high default probability. But implied bond volatilities increase with moneyness (Panel B), and so it turns out that OTM implied bonds have larger Sharpe ratios despite the lower average returns (Panel C).

4.2. One-Month Return Predictability

We forecast each one-month excess return with current option-implied quantities as predictors, pooling all options across strike prices and maturities. For instance, we run the following regression to predict future put option returns with implied volatility:

$$R_{t+1}^{Put}(K_i, T_i) - R_t^f = a + b \times IV_t(K_i, T_i) + \mathbf{c}' \times \text{Controls}_{i,t} + \varepsilon_{i,t+1}$$

We likewise run the same regression with the implied spread $IS_t(K_i, T_i)$ and the normalized implied spread $NIS_t(K_i, T_i)$ as predictors. We repeat the process for hedged put option returns and implied bond excess returns. In each of these regressions, we cluster standard errors by month to correct for the strong cross-sectional correlation among option prices.

Return Predictability Results. Table 2 reports results from pooled predictive regressions. In Panel A are regressions without controls. First, columns 1 to 3 predict the put option excess return in (8), the hedged put option return in (9), and the implied bond excess return in (10) with implied volatility, respectively. Implied volatility predicts neither put option returns (as is well known) nor implied bond returns. Second, columns 4 to 6 predict returns with the implied spread. The coefficient on the implied spread is significant for both put options and implied bonds. But the R^2 s are near zero, which suggests the regression power stems from the large sample size. Finally, columns 7 to 9 predict returns with the normalized implied spread. Like the implied spread, the normalized implied spread is a significant predictor of future excess returns. The lone exception is the coefficient for implied bond returns (column 9), which is only marginally significant.

In Panel B are regressions with controls. The controls are standard option characteristics: moneyness, time-to-maturity, and Greeks. Like the previous regressions, implied volatilities are insignificant and implied spreads are significant for each return. Unlike the previous regressions, normalized implied spreads are insignificant for put options but significant for implied bonds.

In Panel C are the same pooled predictive regressions as in Panel A by moneyness/maturity bin. First, the predictive power of the implied spread and normalized implied spread concentrates among OTM options, which is consistent with the regressions in Panel B. Second, the coefficient on the implied spread is no longer significant for put option returns but remains significant for OTM and ATM implied bond returns. This reinforces the role of the large sample size in Panel A. Third, normalized implied spreads predict future implied bond returns especially well at both short and long horizons.⁷

⁷In the appendix, we show that most of the predictive content of normalized implied spreads comes from

Table 2: One-Month Return Predictability

This table reports estimates from pooled regressions of the form

$$\text{Excess Return}_{i,t+1} = a + b \times \text{Implied Quantity}_{i,t} + \mathbf{c}' \times \text{Controls}_{i,t} + \varepsilon_{i,t+1}$$

The excess return is either the put option return in excess of the risk-free rate, the put option return in excess of a delta-hedged position in the S&P 500 index, or the implied bond return in excess of the risk-free rate. The implied quantity is either the put's implied volatility, implied spread, or normalized implied spread. Controls are the put's moneyness (Mon), time-to-maturity (Dtm), and Greeks (Gamma, Vega, and Theta). OTM is moneyness $K/S \leq 0.95$, ATM is $0.975 \leq K/S \leq 1.025$, and ITM is $K/S \geq 1.05$. Standard errors are clustered by month. The sample is monthly from January 1990 to December 2020.

		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
		Implied Volatility			Implied Spread			Normalized Implied Spread		
		Put	Hedged Put	Implied Bond	Put	Hedged Put	Implied Bond	Put	Hedged Put	Implied Bond
Panel A: Pooled Regressions without Controls										
	b	-0.42	-0.42	0.03	0.14	0.15	0.01	1.57	1.58	0.04
	t	(-1.23)	(-1.29)	(1.37)	(2.29)	(2.60)	(2.85)	(4.37)	(4.57)	(1.75)
	R2	0.00	0.00	0.01	0.00	0.00	0.01	0.01	0.01	0.00
	N	26460	26460	26460	26460	26460	26460	26460	26460	26460
Panel B: Pooled Regressions with Controls										
	b	0.14	0.15	0.03	0.19	0.20	0.02	0.07	0.11	0.08
	t	(0.31)	(0.36)	(1.05)	(2.09)	(2.29)	(2.64)	(0.10)	(0.17)	(1.89)
	Mon	1.07	1.11	0.03	0.72	0.73	0.00	1.03	1.05	0.02
	t	(5.57)	(5.98)	(2.14)	(5.26)	(5.45)	(0.11)	(8.06)	(8.41)	(1.85)
	Dtm	0.10	0.11	0.01	0.15	0.15	0.01	0.11	0.11	-0.00
	t	(1.49)	(1.60)	(1.82)	(1.64)	(1.77)	(2.45)	(1.51)	(1.59)	(-0.01)
	Gamma	-0.05	-0.05	0.00	-0.04	-0.03	0.00	-0.05	-0.05	0.00
	t	(-0.77)	(-0.72)	(1.16)	(-0.52)	(-0.46)	(1.33)	(-0.79)	(-0.74)	(1.30)
	Vega	0.03	0.03	-0.00	0.03	0.03	-0.00	0.02	0.02	-0.00
	t	(3.26)	(3.32)	(-2.16)	(2.72)	(2.78)	(-1.98)	(2.15)	(2.18)	(-2.26)
	Theta	0.07	0.07	-0.00	0.07	0.07	-0.01	0.06	0.06	-0.00
	t	(1.29)	(1.32)	(-1.57)	(1.42)	(1.44)	(-2.02)	(1.22)	(1.24)	(-1.77)
	R2	0.04	0.04	0.02	0.04	0.04	0.02	0.03	0.04	0.02
	N	26460	26460	26460	26460	26460	26460	26460	26460	26460
Panel C: Pooled Regressions by Moneyness and Maturity										
OTM $T \leq 122$	b	0.25	0.26	0.04	0.59	0.60	0.06	3.24	3.26	0.08
	t	(0.50)	(0.53)	(2.56)	(0.94)	(0.98)	(2.53)	(3.18)	(3.27)	(2.38)
	R2	0.00	0.00	0.04	0.00	0.00	0.05	0.01	0.01	0.02
	N	4452	4452	4452	4452	4452	4452	4452	4452	4452
ATM $T \leq 122$	b	0.07	0.08	0.04	-0.20	-0.19	0.04	0.82	0.84	0.08
	t	(0.14)	(0.16)	(1.34)	(-0.82)	(-0.77)	(2.83)	(0.80)	(0.85)	(1.67)
	R2	-0.00	-0.00	0.01	0.00	0.00	0.02	0.00	0.00	0.01
	N	4452	4452	4452	4452	4452	4452	4452	4452	4452
ITM $T \leq 122$	b	-0.25	-0.23	0.05	-0.08	-0.07	0.01	0.01	-0.01	0.04
	t	(-0.60)	(-0.60)	(1.25)	(-1.37)	(-1.35)	(2.50)	(0.01)	(-0.01)	(0.63)
	R2	0.00	0.00	0.01	0.00	0.00	0.00	-0.00	-0.00	0.00
	N	4430	4430	4430	4430	4430	4430	4430	4430	4430
OTM $T \geq 152$	b	-0.11	-0.10	0.05	-0.30	-0.29	0.09	0.16	0.19	0.07
	t	(-0.32)	(-0.29)	(2.19)	(-0.49)	(-0.49)	(2.16)	(0.39)	(0.48)	(3.00)
	R2	-0.00	-0.00	0.03	0.00	0.00	0.03	-0.00	0.00	0.02
	N	4310	4310	4310	4310	4310	4310	4310	4310	4310
ATM $T \geq 152$	b	-0.08	-0.06	0.05	-0.27	-0.26	0.06	0.05	0.07	0.07
	t	(-0.23)	(-0.20)	(1.60)	(-0.71)	(-0.74)	(1.66)	(0.14)	(0.21)	(2.34)
	R2	-0.00	-0.00	0.01	0.00	0.00	0.01	-0.00	-0.00	0.01
	N	4402	4402	4402	4402	4402	4402	4402	4402	4402
ITM $T \geq 152$	b	-0.06	-0.05	0.04	-0.22	-0.21	0.03	0.07	0.07	0.06
	t	(-0.20)	(-0.19)	(1.13)	(-1.39)	(-1.47)	(1.68)	(0.22)	(0.26)	(1.47)
	R2	-0.00	-0.00	0.01	0.00	0.00	0.00	-0.00	-0.00	0.00
	N	4414	4414	4414	4414	4414	4414	4414	4414	4414

Table 3: One-Month Return Predictability: Economic Significance

This table reports the standard deviation of the fitted value (left column) and the ratio of this standard deviation to the unconditional average return (right column) from regressions of one-month excess returns on implied spreads by moneyness and time-to-maturity. The corresponding regressions are in Table 2.

	$\sigma [E_t (R_{t+1}^e)]$		$\frac{\sigma [E_t (R_{t+1}^e)]}{E(R_{t+1}^e)}$	
	Put	Implied Bond	Put	Implied Bond
OTM $T \leq 122$	3.16	0.30	0.09	0.72
OTM $T \geq 152$	0.96	0.29	0.09	0.67

In sum, Table 2 shows that while there is evidence of predictability of implied-bond excess returns, there is little evidence, if any, of predictability of put option excess returns.

To drive home the main point, Table 3 shows that the economic significance of the implied spread is large for implied bonds and small for put options, following the calculation in Cochrane (2011). For implied bonds, the variation in conditional expected returns – the standard deviation of the fitted value in the regression $\sigma [E_t (R_{t+1}^e)] = \sigma (\hat{b} \times IS_t)$ – is more than half as large as unconditional expected returns $E(R_{t+1}^e)$. In contrast, for put options, the variation in conditional expected returns is no more than one-tenth as large as unconditional expected returns. In other words, there is evidence of a substantial time-varying risk premium for implied bonds but not for put options. Similar calculations hold for normalized implied spreads (not tabulated).⁸

Interpreting the Coefficients. Interestingly, the slope coefficient on IV, IS, and NIS are mostly positive for both put options and implied bonds (the only few exceptions are for specifications in which the coefficients are not significant). That both slopes are positive may at first appear puzzling because implied bonds are short put options. In Section 5, we show that a two-factor option pricing model can in fact generate these signs. In this section, we develop simple intuition for these slope coefficients via an identity. By definition, the expected excess return on the implied bond is the weighted average of the expected excess return on the zero-coupon bond and that on the put option:

$$E_t [R_{t+1}^{IB}(K, T) - R_t^f] = \omega_t E_t [R_{t+1}^Z(T) - R_t^f] + (1 - \omega_t) E_t [R_{t+1}^{Put}(K, T) - R_t^f] \quad (11)$$

implied spreads, not implied default frequencies.

⁸Campbell and Thompson (2008) quantify the economic significance of predictive regressions via portfolio choice. If the R^2 is large relative to the squared Sharpe ratio, then a mean-variance investor can use the predictive regression to substantially increase in his single-period portfolio return. For implied bonds, the proportional increase from observing the current implied spread is at least 48% (not tabulated). In contrast, for put options, the proportional increase is at most 1% (not tabulated). This is further evidence the regression R^2 s are economically meaningful for implied bonds but not for put options.

where the weight satisfies

$$\omega_t = \frac{Z_t(T)}{Z_t(T) - P_t(K, T)/K} = \frac{1}{1 - \frac{P_t(K, T)}{Z_t(T)K}} = \exp\{IS_t(K, T)(T - t)\} > 1$$

This identity is the Merton insight – that a corporate bond is economically equivalent to a safe bond and a put option – but in return, not price, units. The positive relation between implied spreads and implied bonds returns is in fact intuitive. Implied bonds and implied spreads are analogous to corporate bonds and corporate bond spreads, respectively. The positive relation follows directly from the strong evidence that corporate bond spreads predict corporate bond returns (see, for example, Nozawa 2017).

As the puzzle is in fact the put option slope, we turn (11) on its head

$$E_t \left[R_{t+1}^{Put}(K, T) - R_t^f \right] = a_t E_t \left[R_{t+1}^{IB}(K, T) - R_t^f \right] + (1 - a_t) E_t \left[R_{t+1}^Z(T) - R_t^f \right] \quad (12)$$

where the weight now satisfies

$$a_t = \frac{1}{1 - \omega_t} = \frac{1}{1 - e^{IS_t(K, T)(T-t)}} < 0 \quad \text{with} \quad \frac{\partial a}{\partial IS} > 0$$

That is, a put option is equivalent to a short position in an implied bond and a long position in a risk-free bond. From (12), when implied spreads increase, there are two effects on expected returns. The direct effect is on the implied bond excess return: higher implied spreads predict higher implied bond returns from the usual risk premium channel. For a given weight a_t , the direct effect implies a negative relation between implied spreads and put option returns because $a_t < 0$. The indirect effect is on the weight a_t – the weight is in fact a function of the current implied spread and so not constant – and higher implied spreads push the weight towards zero. If latter outweighs the former, higher implied spreads would predict higher (less negative) put option returns. The data in Table 2 and stochastic volatility, stochastic jump intensity model in Section 5 support this channel.

4.3. Uncovering Time-Varying Risk Premia: The Joint Return Predictability of Put Options and Implied Bonds

The link between put option and implied bond returns in (11) helps assess statistical significance of our estimates and thus provides further evidence on the time-varying nature of option risk premia. In particular, because the implied spread IS_t is time-varying in the data, the weight $\omega_t = e^{IS_t(T-t)}$ is also time-varying. This implies that *some* predictability of either implied bonds, put options, or the risk-free bond must be in the data. In other words,

denoting by RP^X the risk premium of asset X , we cannot have the identity

$$RP^{IB} = \omega_t RP^Z + (1 - \omega_t) RP^{Put} \quad (13)$$

with RP^{IB} , RP^Z , and RP^{Put} all independent of IS_t , except for knife-edge cases based on their numerical values. In particular, note that if RP^{IB} and RP^{Put} were both unpredictable and independent of IS_t , then RP^Z would have to strongly depend on the option-implied spread IS_t , which is unlikely on an ex-ante basis and in fact finds little support in the data.

More specifically, assuming that the risk premium of zero-coupon Treasury bonds does not depend on the implied spread, then there is no logical null hypothesis in which put options and implied bonds are both unpredictable.⁹ Most important, whether implied spreads predict put option returns or implied bond returns are not separate questions. They are, in fact, the same question. A null hypothesis in which put options are unpredictable, i.e. RP^{Put} in (13) is constant, also specifies that implied bonds are predictable

$$RP_t^{IB} = RP^{Put} + \omega_t (RP^Z - RP^{Put})$$

A null hypothesis in which implied bonds are unpredictable, i.e. RP^{IB} in (13) is constant, also specifies that put options are predictable

$$RP_t^{Put} = RP^{IB} + \frac{\omega_t}{1 - \omega_t} (RP^{IB} - RP^Z)$$

When we evaluate either null, we simultaneously address the joint predictability of put options and implied bonds, whether we acknowledge it or not.

We evaluate joint predictability under each null hypothesis by using Monte Carlo simulations. Denote by b^{Put} and b^{IB} the slope coefficients from predictability regressions of excess returns of put options and implied bonds, respectively, on implied spreads. We begin with the null $b^{Put} = 0$. Consider the representation of implied spreads, zero-coupon bond returns, and put option returns:

$$\begin{aligned} \log IS_{t+1}(K, T) &= \hat{a}^{IS} + \hat{b}^{IS} \log IS_t(K, T) + \varepsilon_{t+1}^{IS} \\ R_{t+1}^Z(T) &= \hat{a}^Z + \hat{b}^Z IS_t(K, T) + \varepsilon_{t+1}^Z \\ R_{t+1}^{Put}(K, T) &= \hat{a}^{Put} + \hat{b}^{Put} IS_t(K, T) + \varepsilon_{t+1}^{Put} \end{aligned} \quad (14)$$

We use the sample estimates of $\hat{b}^{IS}$, $\hat{b}^Z$, and the error covariance matrix. We set $\hat{b}^{Put}$ to the zero null. We initialize the system from the unconditional density:

$$\log IS_0 \sim \mathcal{N} \left(0, (\hat{\sigma}_\varepsilon^{IS})^2 / \left(1 - (\hat{b}^{IS})^2 \right) \right)$$

⁹The argument is reminiscent of analogous argument for stock returns and dividend growth based on the Campbell and Shiller decomposition. If both stock returns and dividend growth are unpredictable, then the price-dividend ratio is constant. Since it is not, there is no logical null hypothesis in which stock returns and dividend growth are both unpredictable. See Cochrane (2008) for more details.

Table 4: Joint Put Option and Implied Bond Return Predictability

The table summarizes the joint distribution of slope coefficients and t-statistics from regressions of one-month put option excess returns and one-month implied bond excess returns on implied spreads in Monte Carlo simulations under the null that put option returns are unpredictable (left column) and the null that implied bond returns are unpredictable (right column). The percentage is the fraction of 10,000 simulations with estimates at least as extreme as the sample estimates by moneyness and time-to-maturity. The corresponding regressions are in Table 2.

	$H_0: b^{Put} = 0$		$H_0: b^{IB} = 0$	
	Slope	t-stat	Slope	t-stat
OTM $T \leq 122$	3.35	8.96	0.00	0.01
OTM $T \geq 152$	28.37	36.33	0.03	2.93

We then draw the shocks ε_t^{IS} , ε_t^Z , and ε_t^{Put} from a multivariate normal distribution and simulate the system forward. We derive the implied bond return via the identity in (11). We produce 10,000 artificial datasets of length $T = 371$ months. In each artificial dataset, we run pooled predictive regressions as in Table 2 by moneyness/maturity bin.

The first null is that put option returns are unpredictable $b^{Put} = 0$. Under this null, implied bond returns are predictable. To evaluate the null, we examine the joint distribution of put option and implied bond coefficients and t-statistics. The left side of Table 4 reports the simulation results. The simulation produces coefficients at least as large as the sample estimates 3% (for short-term options) to 30% (for long-term options) of the time and t-statistics at least 9% of the time. At conventional significance levels, we cannot reject the hypothesis that b^P is zero and b^{IB} is not.

The second null is that implied bond returns are unpredictable $b^{IB} = 0$. Under this null, put option returns are predictable. The simulation is analogous to that in (14), but we replace put options with implied bonds in the system and derive the put option return via the identity in (12). The right side of Table 4 reports the simulation results. The simulation produces coefficients at least as large as the sample estimates less than 1% of the time and t-statistics nearly 0%. At conventional significance levels, we can reject the hypothesis that b^{IB} is zero and b^{Put} is not.

In sum, the joint distribution of put option and implied bond estimates under each null strengthens our results. The failure to reject the null $b^{Put} = 0$ reinforces implied bond returns are predictable. The rejection of the null $b^{IB} = 0$ reinforces our argument that put option returns are unpredictable.¹⁰

¹⁰The evidence that implied spreads predict implied bond returns is subject to the usual statistical biases in predictive regressions. Our sample is relatively short and our predictors extremely persistent. Thus, we

4.4. Long-Horizon Return Predictability

Slow-moving risk premia are easier to detect by running long-horizon return predictability regressions. We begin with implied bonds and then turn to put options.

Implied Bonds. We forecast long-horizon excess returns with current option-implied quantities as predictors by moneyness/maturity bin. For OTM and ATM options with quarterly maturities, we run the following regression to predict future implied bond returns

$$\sum_{i=1}^H \log(1 + R_{t+i}^{IB}(M, H)) - r_{t-1+i}^f = a + \mathbf{b}'\mathbf{X}_t(M, H) + \varepsilon_{t+i}$$

where the left-hand side is the monthly excess return from $t + i - 1$ to $t + i$ of an implied bond with moneyness $M = K/S$ and time-to-maturity H , and the right-hand side predictor is either the implied volatility IV, implied spread IS, or normalized implied spread NIS of the corresponding put option. In other words, the regressions hold constant the moneyness/maturity of the left-hand side portfolio over the horizon.

Table 5 reports results from long-horizon predictive regressions. In Panel A are univariate regressions. With large t-statistics and high R^2 s, both implied spreads and normalized implied spreads strongly predict future implied bond returns. However, as implied default intensities are noisy, normalized implied spreads display lower R^2 s than implied spreads but remain significant. Relative to either implied spreads or normalized implied spreads, implied volatilities are less effective predictors, especially among long-horizon, ATM options.

In Panel B are multivariate regressions. On the one hand, implied spreads always subsume any predictive content from implied volatilities. On the other hand, normalized implied spreads only sometimes subsume implied volatilities. As before, normalized implied spreads contain measurement error, which may diminish their predictive power.

That the Financial Crisis drives the predictability in Table 5 is a source of concern. To address this, we run the same long-horizon regressions in subsamples: January 1990 to June 2005 and July 2005 to December 2020. We report these regressions in the appendix. The coefficients are similar in each sample. In the earlier sample, some coefficients are insignificant due to the lack of power, but the magnitudes are largely consistent with those in the full sample.

Put Options. That implied spreads predict implied bond returns may not be surprising given similar results for corporate bond spreads and corporate bond returns. But the pre-

have a textbook case of Stambaugh bias. However, the estimate of b^{IB} in simulated data accounts for small sample bias, and so the simulations provide further empirical support to implied bond return predictability.

Table 5: Long-Horizon Implied Bond Return Predictability

This table reports estimates from predictive regressions of the form

$$\sum_{i=1}^H \log(1 + R_{t+i}^{IB}(M, H)) - r_{t-1+i}^f = a + \mathbf{b}'\mathbf{X}_t(M, H) + \varepsilon_{t+i}$$

r_{t-1+i}^f is the continuously compounded risk-free rate at $t - 1 + i$. $R_{t+i}^{IB}(M, H)$ is the monthly return from $t + i - 1$ to $t + i$ of an implied bond with moneyness $M = K/S$ and time-to-maturity H . $X_t(M, H)$ is either the implied volatility IV, implied spread IS, or normalized implied spread NIS of a put option with moneyness M and time-to-maturity H . Test statistics are calculated with the Hansen and Hodrick standard error correction for overlapping data. The sample is monthly from January 1990 to December 2020.

Panel A: Univariate Regressions															
		Implied Volatility				Implied Spread				Normalized Implied Spread					
		Horizon (Months)				Horizon (Months)				Horizon (Months)					
K/S		3	6	9	12	3	6	9	12	3	6	9	12		
0.900	b	0.13	0.27	0.37	0.41	0.19	0.51	0.81	1.02	0.31	0.45	0.48	0.49		
	t	(3.73)	(5.24)	(4.19)	(4.12)	(2.31)	(3.89)	(4.37)	(5.58)	(4.94)	(4.90)	(3.66)	(2.95)		
	R2	0.15	0.19	0.17	0.15	0.13	0.19	0.19	0.19	0.12	0.15	0.11	0.09		
	N	366	348	322	294	366	348	322	294	366	348	322	294		
0.950	b	0.14	0.27	0.35	0.42	0.17	0.46	0.67	0.98	0.39	0.49	0.52	0.57		
	t	(2.67)	(3.38)	(3.02)	(2.99)	(2.11)	(3.56)	(3.39)	(4.01)	(3.86)	(4.19)	(3.34)	(2.98)		
	R2	0.09	0.11	0.10	0.10	0.08	0.12	0.10	0.12	0.09	0.10	0.09	0.08		
	N	366	360	351	343	366	360	351	343	366	360	351	343		
0.975	b	0.14	0.26	0.33	0.39	0.19	0.45	0.75	0.98	0.41	0.46	0.52	0.56		
	t	(2.28)	(2.97)	(2.48)	(2.41)	(2.44)	(3.16)	(3.42)	(3.62)	(3.47)	(3.26)	(2.75)	(2.68)		
	R2	0.07	0.08	0.07	0.07	0.08	0.10	0.11	0.11	0.07	0.07	0.07	0.07		
	N	366	360	354	348	366	360	354	348	366	360	354	348		
1.000	b	0.14	0.24	0.32	0.37	0.15	0.41	0.72	0.94	0.32	0.50	0.55	0.60		
	t	(1.78)	(2.14)	(2.06)	(2.02)	(1.72)	(2.94)	(2.83)	(3.42)	(2.15)	(2.93)	(2.62)	(2.70)		
	R2	0.04	0.05	0.06	0.05	0.04	0.08	0.09	0.09	0.03	0.06	0.06	0.06		
	N	366	366	354	348	366	366	354	348	366	366	354	348		
Panel B: Multivariate Regressions															
		IS and IV				NIS and IV				IS and NIS					
		Horizon (Months)				Horizon (Months)				Horizon (Months)					
K/S		3	6	9	12	3	6	9	12	3	6	9	12		
0.900	IS	0.02	0.33	0.67	1.31	NIS	0.08	0.12	-0.00	-0.08	IS	0.13	0.41	0.79	1.24
	t	(0.13)	(1.01)	(2.05)	(2.86)	t	(1.10)	(0.89)	(-0.01)	(-0.32)	t	(0.99)	(1.86)	(3.03)	(5.05)
	IV	0.12	0.11	0.07	-0.15	IV	0.11	0.22	0.37	0.46	NIS	0.15	0.14	0.02	-0.20
	t	(1.46)	(0.89)	(0.60)	(-0.64)	t	(2.09)	(2.46)	(2.48)	(2.78)	t	(1.20)	(0.95)	(0.18)	(-0.80)
	R2	0.15	0.20	0.19	0.19	R2	0.15	0.19	0.17	0.15	R2	0.14	0.20	0.19	0.20
	N	366	348	322	294	N	366	348	322	294	N	366	348	322	294
0.950	IS	-0.05	0.43	0.43	1.12	NIS	0.20	0.19	0.19	0.15	IS	0.07	0.35	0.47	0.93
	t	(-0.26)	(1.53)	(1.23)	(1.48)	t	(1.27)	(0.83)	(0.68)	(0.46)	t	(0.54)	(1.61)	(1.82)	(3.72)
	IV	0.18	0.02	0.14	-0.08	IV	0.08	0.19	0.25	0.34	NIS	0.26	0.17	0.22	0.05
	t	(1.85)	(0.11)	(0.66)	(-0.20)	t	(0.92)	(1.21)	(1.08)	(1.25)	t	(1.51)	(0.93)	(1.44)	(0.30)
	R2	0.09	0.12	0.10	0.12	R2	0.09	0.11	0.10	0.10	R2	0.09	0.12	0.10	0.12
	N	366	360	351	343	N	366	360	351	343	N	366	360	351	343
0.975	IS	0.28	0.47	1.18	1.49	NIS	0.27	0.13	0.24	0.27	IS	0.12	0.38	0.69	0.99
	t	(2.37)	(1.67)	(2.34)	(1.87)	t	(1.48)	(0.58)	(0.76)	(0.70)	t	(1.06)	(1.87)	(2.76)	(3.45)
	IV	-0.08	-0.01	-0.25	-0.28	IV	0.06	0.20	0.21	0.23	NIS	0.20	0.10	0.06	-0.01
	t	(-0.92)	(-0.09)	(-0.91)	(-0.68)	t	(0.61)	(1.36)	(0.83)	(0.76)	t	(1.34)	(0.52)	(0.36)	(-0.04)
	R2	0.08	0.09	0.11	0.11	R2	0.07	0.08	0.08	0.07	R2	0.08	0.10	0.11	0.11
	N	366	360	354	348	N	366	360	354	348	N	366	360	354	348
1.000	IS	0.07	0.54	0.89	1.25	NIS	0.13	0.34	0.36	0.45	IS	0.11	0.31	0.60	0.86
	t	(0.55)	(3.54)	(1.97)	(1.72)	t	(0.83)	(1.45)	(1.05)	(1.08)	t	(1.05)	(1.94)	(1.80)	(2.54)
	IV	0.08	-0.10	-0.11	-0.18	IV	0.10	0.11	0.16	0.13	NIS	0.14	0.19	0.14	0.08
	t	(0.75)	(-0.73)	(-0.42)	(-0.46)	t	(1.07)	(0.67)	(0.61)	(0.38)	t	(0.96)	(1.04)	(0.63)	(0.33)
	R2	0.04	0.08	0.08	0.10	R2	0.04	0.06	0.07	0.06	R2	0.04	0.08	0.09	0.09
	N	366	366	354	348	N	366	366	354	348	N	366	366	354	348

dictability of implied bonds stands in stark contrast to that of put options. Table 6 reports long-horizon regressions for put option returns, as in Table 5 for implied bond returns. Neither IV (as is well known), IS, nor NIS predict future put option returns. No coefficient is significant, and the coefficient signs are not consistent across strike prices or horizons. Therefore, put option returns are not predictable.

In a recent paper, Israelov and Kelly (2017) propose an elaborate Monte Carlo simulation to forecast the distribution of option returns. Their methodology successfully predicts very short-horizon option returns (at the one-day, one-week, and two-week horizons).¹¹ But long-horizon option return predictability is difficult, if not impossible, to pull off.

The dichotomy between strong implied bond return predictability in Table 5 and weak put option return predictability in Table 6 is intriguing. We next turn to two dynamic option pricing models to rationalize these results.

5. Risk Premia in Dynamic Option Pricing Models

In this section, we build a model that rationalizes the empirical patterns we documented in the previous sections. Section 2.4 examines option-implied spreads in log-normal and jump diffusion models. It is difficult, however, to assess return predictability in these models because neither has a dynamic state variable. This section instead considers two dynamic models: the first is a model with stochastic volatility and jumps and the second is a model where jumps also arrive with stochastic intensity.

5.1. Stochastic Volatility Jump Model (SVJ)

Dynamics. We begin with the simpler, stochastic volatility jump (SVJ) model to discuss its virtues and its shortcomings. The SVJ model is a standard, one-factor model. Specifically, the stock return follows a jump-diffusion process and volatility follows a square-root process

$$dS_t/S_t = [\mu_S - \delta - \lambda E(J_S - 1)] dt + \sqrt{v_t} dW_{S,t} + (J_S - 1) dQ_t \quad (15)$$

$$dv_t = \kappa_v(\theta_v - v_t) dt + \sigma_v \sqrt{v_t} dW_{v,t} \quad (16)$$

where $\rho = E[W_{S,t}W_{v,t}]$ is the correlation between return and volatility shocks, J_S , a random variable with $\log(1 + J_S) \sim \mathcal{N}(\log(1 + \mu_J) - \frac{1}{2}\sigma_J^2, \sigma_J^2)$, determines the stochastic jump size,

¹¹Because conventional returns on deep OTM options are often extreme, Israelov and Kelly predict delta-hedged option returns relative to the initial underlying stock price, not the initial option price.

Table 6: Long-Horizon Put Option Return Predictability

This table reports estimates from predictive regressions of the form

$$\sum_{i=1}^H \log(1 + R_{t+i}^{Put}(M, H)) - r_{t-1+i}^f = a + \mathbf{b}'\mathbf{X}_t(M, H) + \varepsilon_{t+i}$$

r_{t-1+i}^f is the continuously compounded risk-free rate at $t - 1 + i$. $R_{t+i}^{Put}(M, H)$ is the monthly return from $t + i - 1$ to $t + i$ of a put option with moneyness $M = K/S$ and time-to-maturity H . $X_t(M, H)$ is either the implied volatility IV, implied spread IS, or normalized implied spread NIS of a put option with moneyness M and time-to-maturity H . Test statistics are calculated with the Hansen and Hodrick standard error correction for overlapping data. The sample is monthly from January 1990 to December 2020.

Panel A: Univariate Regressions													
K/S		Implied Volatility				Implied Spread				Normalized Implied Spread			
		Horizon (Months)				Horizon (Months)				Horizon (Months)			
		3	6	9	12	3	6	9	12	3	6	9	12
0.900	b	-0.45	-0.29	-0.77	-0.78	-0.29	-1.85	-3.11	-3.66	2.73	0.46	0.32	1.04
	t	(-0.27)	(-0.19)	(-0.42)	(-0.37)	(-0.10)	(-0.54)	(-0.82)	(-0.75)	(0.80)	(0.18)	(0.13)	(0.31)
	R2	-0.00	-0.00	-0.00	-0.00	-0.00	-0.00	0.00	0.00	-0.00	-0.00	-0.00	-0.00
	N	366	348	322	294	366	348	322	294	366	348	322	294
0.950	b	-0.54	-0.27	-0.24	-0.35	-0.20	-1.09	-0.42	-2.65	0.51	-0.15	-0.68	-0.78
	t	(-0.31)	(-0.18)	(-0.12)	(-0.17)	(-0.09)	(-0.43)	(-0.11)	(-0.61)	(0.14)	(-0.06)	(-0.27)	(-0.29)
	R2	-0.00	-0.00	-0.00	-0.00	-0.00	-0.00	-0.00	0.00	-0.00	-0.00	-0.00	-0.00
	N	366	360	351	343	366	360	351	343	366	360	351	343
0.975	b	-0.26	-0.37	0.11	-0.32	-0.64	-1.17	-1.88	-3.48	-0.46	-0.13	-0.52	-0.96
	t	(-0.16)	(-0.23)	(0.06)	(-0.15)	(-0.33)	(-0.46)	(-0.53)	(-0.80)	(-0.14)	(-0.05)	(-0.20)	(-0.33)
	R2	-0.00	-0.00	-0.00	-0.00	-0.00	-0.00	0.00	0.01	-0.00	-0.00	-0.00	-0.00
	N	366	360	354	348	366	360	354	348	366	360	354	348
1.000	b	-0.89	-0.16	0.03	-0.10	-0.77	-1.47	-1.75	-3.25	-0.65	-1.12	-1.24	-1.31
	t	(-0.54)	(-0.10)	(0.02)	(-0.05)	(-0.41)	(-0.70)	(-0.49)	(-0.83)	(-0.19)	(-0.43)	(-0.47)	(-0.45)
	R2	-0.00	-0.00	-0.00	-0.00	-0.00	0.00	0.00	0.01	-0.00	-0.00	-0.00	-0.00
	N	366	366	354	348	366	366	354	348	366	366	354	348
Panel B: Multivariate Regressions													
K/S		IS and IV				NIS and IV				IS and NIS			
		Horizon (Months)				Horizon (Months)				Horizon (Months)			
		3	6	9	12	3	6	9	12	3	6	9	12
0.900	IS	2.09	-8.23	-10.71	-14.14	NIS	10.31	2.60	3.79	5.99	IS	-3.42	-5.13
	t	(0.30)	(-1.04)	(-1.72)	(-1.41)	t	(2.10)	(0.53)	(0.83)	(0.98)	t	(-0.78)	(-0.91)
	IV	-1.72	3.70	3.97	5.27	IV	-3.58	-1.39	-2.68	-3.97	NIS	7.26	4.38
	t	(-0.45)	(1.15)	(1.52)	(1.41)	t	(-1.45)	(-0.48)	(-0.82)	(-1.09)	t	(1.63)	(1.03)
0.950	R2	-0.00	0.00	0.01	0.01	R2	0.01	-0.00	0.00	0.01	R2	0.00	0.00
	N	366	348	322	294	N	366	348	322	294	N	366	348
	IS	5.43	-6.93	-0.04	-12.84	NIS	6.09	1.03	-1.27	-1.20	IS	-1.12	-2.60
	t	(1.05)	(-1.03)	(-0.00)	(-0.88)	t	(0.94)	(0.19)	(-0.23)	(-0.20)	t	(-0.28)	(-0.61)
0.975	IV	-4.70	3.78	-0.22	5.35	IV	-2.40	-0.73	0.45	0.34	NIS	2.51	2.27
	t	(-1.46)	(0.91)	(-0.05)	(0.78)	t	(-0.76)	(-0.23)	(0.10)	(0.07)	t	(0.42)	(0.54)
	R2	-0.00	0.00	-0.01	0.01	R2	-0.00	-0.01	-0.01	-0.01	R2	-0.00	-0.00
	N	366	360	351	343	N	366	360	351	343	N	366	360
1.000	IS	-4.69	-5.53	-14.05	-16.71	NIS	0.36	1.48	-2.14	-1.95	IS	-1.31	-2.82
	t	(-1.25)	(-0.95)	(-1.53)	(-1.25)	t	(0.07)	(0.30)	(-0.38)	(-0.31)	t	(-0.47)	(-0.75)
	IV	3.55	2.87	7.09	7.10	IV	-0.37	-0.99	1.24	0.79	NIS	1.91	2.56
	t	(1.10)	(0.83)	(1.44)	(1.11)	t	(-0.14)	(-0.34)	(0.30)	(0.17)	t	(0.43)	(0.70)
1.000	R2	-0.00	-0.00	0.02	0.03	R2	-0.01	-0.00	-0.00	-0.00	R2	-0.00	-0.00
	N	366	360	354	348	N	366	360	354	348	N	366	360
	IS	1.56	-5.99	-7.97	-11.82	NIS	2.32	-2.07	-3.04	-3.40	IS	-1.18	-1.77
	t	(0.44)	(-1.97)	(-1.08)	(-1.18)	t	(0.56)	(-0.49)	(-0.59)	(-0.59)	t	(-0.53)	(-0.68)
1.000	IV	-2.18	3.61	3.83	5.15	IV	-1.53	0.64	1.42	1.72	NIS	1.31	0.60
	t	(-0.72)	(1.59)	(1.02)	(0.98)	t	(-0.73)	(0.26)	(0.41)	(0.40)	t	(0.37)	(0.19)
	R2	-0.00	0.01	0.01	0.02	R2	-0.00	-0.00	-0.00	-0.00	R2	-0.00	-0.00
	N	366	366	354	348	N	366	366	354	348	N	366	366

Table 7: SVJ Parameter Estimates

The table reports parameter estimates for the SVJ model from Broadie, Chernov, and Johannes (2009) and Chambers, Foy, Liebner, and Lu (2014).

r	δ	μ_S	λ^P	μ_J^P	σ_J^P	κ^P	θ^P	σ_v	$-\rho$
3.70%	2.00%	5.83%	1.29	-1.22%	8.94%	2.192	0.0171	0.258	-0.760
			λ^Q	μ_J^Q	σ_J^Q	ξ_v			
			1.29	-1.22%	8.94%	0			

and dQ_t , the increment of a Poisson process with constant intensity λ , determines jump arrivals. We relax the assumption of constant intensity when we turn to the stochastic intensity model next.

The state price density M_t follows the process

$$dM_t/M_t = [-r - \lambda E(J_M - 1)] dt - \sigma_{M,S}(v_t) dW_{S,t} - \xi_v \sqrt{v_t} dW_{v,t} + (J_M - 1) dQ_t$$

where $\xi_v \sqrt{v_t}$ is the market price of volatility risk. Standard arguments imply that the risk premium of any security $V(S, v, t)$ is

$$E_t \left[\frac{dV_t}{V_t} \right] / dt - r_t = \mu_{V,t} = -Cov^{diff} \left(\frac{dV_t}{V_t}, \frac{dM_t}{M_t} \right) / dt - Cov^{jump} \left(\frac{dV_t}{V_t}, \frac{dM_t}{M_t} \right) / dt \quad (17)$$

$$= \beta_S^V \mu_S + \beta_v^V \xi_v \sigma_v v_t - \lambda E_t [(J_V - 1)(J_M - 1)] \quad (18)$$

where $\beta_S^V = \frac{\partial \log V}{\partial \log S}$ is the loading on stock risk, $\mu_S = \sigma_{M,S}(v_t) \sqrt{v_t}$ is the stock risk premium, and $\beta_v^V = \frac{\partial \log V}{\partial v}$ is the loading on volatility risk. As in Broadie, Chernov, and Johannes (2009) we assume μ_S constant and independent of stock volatility.


Table 7 reports baseline parameter estimates from Broadie, Chernov, and Johannes (2009) and Chambers, Foy, Liebner, and Lu (2014) with one exception. These papers estimate the model with futures options and set the dividend yield δ to the risk-free rate. Since we are interested in predictions for the spot market, we set $\delta = 2\%$.

Monte Carlo Simulations. Figure 10 plots the relation between option-implied quantities and risk premia from Monte Carlo simulations of the SVJ model. We simulate the model under the physical measure P and price options under the risk-neutral measure Q .

Consider first the put options on the left panels of Figure 10. There is a positive association between implied volatilities (top panel), implied spreads (middle panel), and normalized implied spreads (bottom panel) and put risk premia. The main reason is the well-known fact that the put option beta $\beta_{S,t}^V = \frac{\partial \log V}{\partial \log S}$, which is negative, becomes less negative for higher volatility v_t . Intuitively, a high volatility of stock returns decreases the impact of the

Figure 10: Expected Excess Returns in SVJ

This figure plots one-month expected excess returns on put options (left column) and implied bonds (right column) by implied volatilities (top row), implied spreads (middle row), and normalized implied spreads (bottom row) from Monte Carlo simulations of the SVJ model. Options have 365 days-to-maturity.



current stock price on the value of the option, as a high volatility makes it similarly likely for the stock to end up in-the-money at maturity for similarly situated stock prices. We note that in Table 7 the market price of volatility risk $\xi_v = 0$ as in Broadie, Chernov, and Johannes (2009) and Chambers, Foy, Liebner, and Lu (2014). A large negative value of ξ_v can potentially overturn this result, as $\beta_{v,t}^V = \frac{\partial \log V}{\partial v} > 0$. However, only a knife-edge value would make the pattern flat, as suggested by the lack of predictability of option returns documented in Sections 4.3 and 4.4 (see the appendix for robustness checks with respect to the market prices of risk).

We next examine implied bonds on the right of Figure 10. Even in this case, we find a positive association between implied volatilities, implied spreads, normalized implied spreads and implied bond risk premia. The positive slope is smaller for ATM and ITM bonds. That is, the SVJ model is qualitatively consistent with implied bond return predictability in the data. Economically, the risk premium of implied bonds can be written as

$$\mu_{B,t} = \omega_t \mu_Z + (1 - \omega_t) \mu_{P,t}$$

where $\omega_t = (1 - P_t(K, T)/Z_t(T))^{-1} > 1$ is increasing in volatility $\sqrt{v_t}$, μ_Z is the risk premium on the zero-coupon bond, and $\mu_{P,t}$ is the risk premium on the put option (see discussion in Section 4.2). On the one hand, higher volatility $\sqrt{v_t}$ implies a less negative $\mu_{P,t}$, which decreases $\mu_{B,t}$ for a given ω_t . On the other hand, higher volatility $\sqrt{v_t}$ increases ω_t and hence decreases the weight on $\mu_{P,t} < 0$ and gives more weight to $\mu_{Z,t} > 0$. Under the parameters in Table 7 the second channel dominates. Note that a negative market price of volatility risk ξ_v in this case would reinforce the predictability of bond returns, as $\beta_{v,t}^V = \frac{\partial \log V}{\partial v} < 0$ for implied bonds.

In sum, the SVJ model is consistent with the predictability of implied bond returns, but it also implies predictability of option returns. As is well known for IV and as we show for IS/NIS, this relation is not in the data. As evidence, none of the put option coefficients in binned regressions (Panel C of Table 2) or long-horizon regressions (Panel A of Table 6) are significant. Neither IV, IS, nor NIS predicts put option returns. Although the relation may be nonlinear – some coefficients are significant in multivariate regressions (Panel B of Table 2) – put option risk premia are generally not in line with the simple SVJ model.

5.2. Stochastic Volatility, Stochastic Jump Intensity Model (SVSIJ)

The counterfactual predictions of the SVJ model suggest a two-factor model to rationalize theory with data. In particular, the model needs to simultaneously generate predictable implied bond returns and unpredictable put option returns. We examine here a two-factor, stochastic volatility, stochastic jump intensity model (SVSIJ).

Dynamics. Under SVSIJ, the stock price, volatility, and jump intensity follow the processes

$$d \log S_t = \left[\mu_S - \delta - \frac{1}{2}v_t - \lambda_t E(J_S - 1) \right] dt + \sqrt{v_t} dW_{S,t} + J_S dQ_t^{\lambda_t} \quad (19)$$

$$dv_t = k_v (\theta_v - v_t) dt + \sigma_v \sqrt{v_t} dW_{v,t} \quad (20)$$

$$d\lambda_t = k_\lambda (\theta_\lambda - \lambda_t) dt + \sigma_\lambda \sqrt{\lambda_t} dW_{\lambda,t} \quad (21)$$

where $\rho = E[W_{S,t}W_{v,t}]$ is the correlation between return and volatility shocks, J_S , a random variable with $J_S \sim \mathcal{N}(\mu_J, \sigma_J^2)$, determines the stochastic jump size, and $dQ_t^{\lambda_t}$, the increment of a Poisson process with intensity λ_t , determines jump arrivals.

Like SVJ, there is stochastic volatility (20) and jumps in returns. But unlike SVJ, jumps arrive with stochastic intensity (21). In particular, volatility and intensity follow independent stochastic processes as $E[W_{S,t}W_{\lambda,t}] = E[W_{v,t}W_{\lambda,t}] = 0$. In contrast, many (but not all) existing models with stochastic intensity parametrize intensity as a deterministic function of

Table 8: SVSIJ Parameter Estimates

This table reports parameter estimates for the SVSIJ model.

r	δ	μ_S	μ_J^P	σ_J^P	κ_v	θ_v	σ_v	ξ_v	ρ
3.70%	2.00%	5.83%	-4.01%	8.94%	1.540	0.0192	0.2433	-3.79	-0.8510
			μ_J^Q	σ_J^Q	κ_λ	θ_λ	σ_λ	ξ_λ	
			-6.58%	8.94%	1.562	0.3412	0.3620	-2.01	

volatility. In other words, our model is indeed a two-factor model and not simply a one-factor model with stochastic volatilities, jumps, and stochastic jump intensities.

The state price density M_t follows the process

$$dM_t/M_t = [-r - \lambda_t E(J_M - 1)] dt - \sigma_{M,S} dW_{S,t} - \xi_v \sqrt{v_t} dW_{v,t} - \xi_\lambda \sqrt{\lambda_t} dW_{\lambda,t} + (J_M - 1) dQ_t^{\lambda_t}$$

where $\xi_v \sqrt{v_t}$ is the market price of volatility risk and $\xi_\lambda \sqrt{\lambda_t}$ is that of intensity risk. Then the risk premium of any security $V(S, v, \lambda, t)$ is

$$E_t \left[\frac{dV_t}{V_t} \right] / dt - r_t = \mu_{V,t} = -Cov^{diff} \left(\frac{dV_t}{V_t}, \frac{dM_t}{M_t} \right) / dt - Cov^{jump} \left(\frac{dV_t}{V_t}, \frac{dM_t}{M_t} \right) / dt \quad (22)$$

$$= \beta_S^V \mu_S + \beta_v^V \xi_v \sigma_v v_t + \beta_\lambda^V \xi_\lambda \sigma_\lambda \lambda_t - \lambda_t E_t [(J_V - 1)(J_M - 1)] \quad (23)$$


where $\beta_S^V = \frac{\partial \log V}{\partial \log S}$ is the loading on stock risk, $\beta_v^V = \frac{\partial \log V}{\partial v}$ on volatility risk, $\beta_\lambda^V = \frac{\partial \log V}{\partial \lambda}$ on intensity risk. Although both volatility and intensity increase IV, IS, and NIS, they may have opposite effects on put option risk premia. If put risk premia increase with volatility (as in SVJ) but decrease with intensity, then put option returns may not be predictable under SVSIJ.

Table 8 reports parameter estimates for the SVSIJ model. We estimate parameters via maximum likelihood (estimation) subject to a return predictability constraint (calibration). The constraint incorporates implied bond and put option return predictability (moment conditions) into the estimation in the spirit of indirect inference (Gourieroux, Monfort, and Renault 1993) and simulated method of moments (Duffie and Singleton 1993). We relegate further discussion of the estimation, calibration, and parameters to the appendix (see Figures D2 to D4 and Tables D6 to D7) and focus on return predictability in this section. We begin with the qualitative predictions from Monte Carlo simulations and then turn to quantitative predictions from simulated data.

Monte Carlo Simulations. Figure 11 plots option-implied spreads and risk premia by volatility and intensity from Monte Carlo simulations of the SVSIJ model. As before with SVJ, we simulate the model under the physical measure P and price options under the risk-neutral measure Q . To evaluate the predictions of the SVSIJ model, we examine a single

Figure 11: Expected Excess Returns by Volatility/Intensity in SVSIJ

This figure plots implied spreads (Panel A), normalized implied spreads (Panel B), one-month expected excess returns on put options (Panel C), and one-month expected excess returns on implied bonds (Panel D) by volatility and intensity from Monte Carlo simulations of the SVSIJ model. Options are out-of-the-money with 365 days-to-maturity.



OTM option with 365 days-to-maturity. We show all predictions hold for shorter-maturity options in the appendix.

In Panel A are implied spreads, Panel B normalized implied spreads, and Panel D implied bond risk premia. Each increases with volatility and intensity. Implied spreads increase because put option prices rise. Normalized implied spreads rise because the increase in implied spreads outweighs that of implied default intensities. Implied bond risk premia increase because default probabilities rise and default carries a risk premium. As spreads and risk premia move together, IS and NIS unsurprisingly predict implied bond returns.

In Panel C are put option risk premia. As with SVJ in Figure 10, put option risk premia increase with volatility. But Panel C shows that risk premia decrease with intensity at low volatilities and are flat at high volatilities. The intuition is that put options hedge stock market jumps. As the last term $\lambda_t E_t[(J_V - 1)(J_M - 1)]$ in (23) is positive, these hedging services command a negative risk premium. As the jump intensity λ_t increases, the hedge becomes more valuable, the hedge risk premium becomes more negative, and so the total risk premium decreases. The estimated mean jump size is sufficiently negative and the market price of intensity risk ξ_λ is sufficiently small that the decrease in the hedge risk premium

Table 9: One-Month Return Predictability in SVSIJ

This table reports estimates from predictive regressions of one-month put option excess returns and one-month implied bond excess returns on implied spreads (left column) and normalized implied spreads (right column) in simulations of the SVSIJ model. The standalone value is the approximate population value in a single, long sample. The brackets contain the 5th and 95th percentile from 100,000 samples of 371 months. Options are out-of-the-money with 365 days-to-maturity. Table 2 reports analogous regressions in the data.

	Implied Spread		Normalized Implied Spread	
	Put	Implied Bond	Put	Implied Bond
b	0.02	0.06	0.02	0.04
	[-0.08, 0.11]	[0.03, 0.10]	[-1.38, 1.39]	[0.02, 0.07]
R2	0.00	0.03	0.00	0.02
	[-0.00, 0.01]	[0.00, 0.06]	[-0.00, 0.01]	[0.00, 0.05]

overpowers any potential increase in the diffusive intensity premium $\beta_\lambda^V \xi_\lambda \sigma_\lambda \lambda_t$. As spreads and risk premia do not necessarily move together, IS and NIS – and IV, although not shown – may not predict put option returns.

Simulated Data. To quantitatively validate the SVSIJ model, Table 9 reports results from predictive regressions in simulated data. The standalone value is the population value from a single long sample. The brackets contain the 5th and 95th percentiles from many short samples. We again relegate details to the appendix and focus on return predictability in this section.

First, on the left are regressions with implied spreads. While the slope for implied bond returns is positive in small samples, the slope is smaller than that in the data (0.15, not tabulated). In contrast, the slope for put option returns is not significant in small samples, as in the data. Second, on the right are regressions with normalized implied spreads. Again, the implied bond slope is significant in small samples, but the put option slope is not. Third, the R^2 s for implied bonds dwarf those of put options. But the 5th percentile in regressions of implied bond returns on normalized implied spreads is 0%, which is again smaller than that in the data (2%, not tabulated). While the quantitative differences in moments is evidence of either imprecise regression estimates or model misspecification, the simulated data are generally consistent with the data.

In short, a two-factor SVSIJ model matches the predictability of implied bond returns and the lack of predictability of put option returns as evident in the data. Under SVSIJ, volatility and intensity have opposite effects on put option risk premia, and so put option returns are not predictable.

Of course, we are not the first to propose a two-factor model with time-varying volatility and time-varying jump intensity for derivatives pricing. In the context of index options, Santa-Clara and Yan (2010) and Christoffersen, Jacobs, and Ornthanalai (2012) show that a two-factor model outperforms (in terms of fit) a standard one-factor model. In the context of singlename credit default swaps, Kelly, Manzo, and Palhares (2019) show that the cross-sectional dynamics of spreads across firms and maturities requires a two-factor model. Our contribution is to show that time-varying jump intensity resolves the shortcomings of the SVJ model in explaining the puzzling dynamics of option risk premia. In sum, implied spreads (or implied volatility) can increase for two reasons, either return volatility increases or jump-intensity increases. Higher volatility makes put option premia higher (less negative) while higher intensity makes put option premia lower (more negative) because of puts' hedging properties. Thus, IV or IS/NIS do not predict option returns. In contrast, because put option risk premia are negative, higher implied spreads always decreases the weight given to put option returns in the implied bond (see (11)) and thus always increase its risk premium.

6. Conclusions

We propose option-implied bonds and the related option-implied spreads and normalized implied spreads as informative measures to uncover the hidden properties of option price dynamics and their risk premia. The implied spread is the credit spread implicit in an option-implied bond, a simple portfolio long a risk-free bond and short a put option and economically equivalent to a corporate bond. The normalized implied spread scales the implied spread by the bond's implicit default intensity. We can readily compute all of these quantities from option prices and risk-free rates alone.

The surface of (normalized) implied spreads is countercyclical: spreads are high in bad times and low in good times. In particular, the shape of the normalized implied spread surface captures the tails of the underlying stock (risk-neutral) distribution. The curvature is instead procyclical: the normalized implied spread surface is convex in good times and concave in bad times. That is, jump risk is more important during good times, as the likelihood of a market crash is high and volatility is low.

Just as corporate bond spreads predict corporate bond returns, we find that implied spreads strongly predict implied bond returns. High implied spreads reliably precede high implied bond returns. Low spreads precede low returns. Implied bond predictability persists up to one-year into the future. In contrast, implied spreads do not predict put option returns. Option returns, especially at long horizons, are difficult to predict.

We find that a standard, one-factor, stochastic volatility jump model cannot simultaneously match implied bond and put option return predictability. We propose an alternative, two-factor, stochastic volatility, stochastic jump intensity model and show that it matches these return predictability moments. Our examination of option-implied spreads and return predictability of both implied bond returns and option returns thus brings forth the necessity of two factors – time-varying volatility and time-varying jump intensity – to describe the elusive nature of option risk premia.

In this paper, we only focus on S&P 500 options in order to study the time-variation in option risk premia at the aggregate level, i.e. securities affected by systematic jump risk and systematic volatility risk. Future research may explore a similar strategy – namely, the study of the joint predictability of implied bond returns and option returns – to learn about the risk premia of options on individual stocks, commodities, or exchange rates. This research may explore the dynamics of risk premia both in the time-series and in the cross-section. While additional challenges present themselves in the study of individual stock options and the definition of implied bonds, we believe such research may yield quite interesting results.

Future research may also explore what preference-based models are able to replicate the predictability of implied bonds and lack of predictability of put option returns. Such research may provide further insights on the type of parameter restrictions on preferences and the underlying dynamics of fundamentals.

REFERENCES

- David S. Bates. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics*, 94(1-2):181–238, 2000.
- David S. Bates. Maximum Likelihood Estimation of Latent Affine Processes. *Review of Financial Studies*, 19(3):909–965, 2006.
- David S. Bates. U.S. stock market crash risk, 1926–2010. *Journal of Financial Economics*, 105(2):229–259, 2012.
- Fischer Black and Myron Scholes. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
- Mark Broadie, Mikhail Chernov, and Michael Johannes. Understanding Index Option Returns. *Review of Financial Studies*, 22(11):4493–4529, 2009.
- John Y. Campbell and Robert J. Shiller. The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. *Review of Financial Studies*, 1(3):195–228, 1988.
- John Y. Campbell and Samuel B. Thompson. Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average? *Review of Financial Studies*, 21(4):1509–1531, 2008.
- Jie Cao, Amit Goyal, Xiao Xiao, and Xintong Zhan. Implied Volatility Changes and Corporate Bond Returns. *Swiss Finance Institute Research Paper No. 19-75*, 2019.
- Donald R. Chambers, Matthew Foy, Jeffrey Liebner, and Qin Lu. Index Option Returns: Still Puzzling. *Review of Financial Studies*, 27(6):1915–1928, 2014.
- Peter Christoffersen, Kris Jacobs, and Chayawat Ornthanalai. Dynamic jump intensities and risk premiums: Evidence from S&P500 returns and options. *Journal of Financial Economics*, 106(3):447–472, 2012.
- Peter Christoffersen, Kris Jacobs, and Bo Young Chang. Forecasting with Option-Implied Information. *Handbook of Economic Forecasting*, 2:581–656, 2013.
- John H. Cochrane. The Dog That Did Not Bark: A Defense of Return Predictability. *Review of Financial Studies*, 21(4):1533–1575, 2008.
- John H. Cochrane. Presidential Address: Discount Rates. *Journal of Finance*, 66(4):1047–1108, 2011.
- George M. Constantinides, Jens Carsten Jackwerth, and Alexi Savov. The Puzzle of Index Option Returns. *Review of Asset Pricing Studies*, 3(2):229–257, 2013.
- Christopher L. Culp, Yoshio Nozawa, and Pietro Veronesi. Option-Based Credit Spreads. *American Economic Review*, 108(2):454–88, 2018.

- Alexander David and Pietro Veronesi. Investors' and Central Bank's Uncertainty Embedded in Index Options. *Review of Financial Studies*, 27(6):1661–1716, 2014.
- Darrell Duffie and Kenneth Singleton. Simulated Moments Estimation of Markov Models of Asset Prices. *Econometrica*, pages 929–952, 1993.
- Darrell Duffie, Jun Pan, and Kenneth Singleton. Transform Analysis and Asset Pricing for Affine Jump-Diffusions. *Econometrica*, 68(6):1343–1376, 2000.
- Bjørn Eraker. Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices. *Journal of Finance*, 59(3):1367–1403, 2004.
- Christian Gourieroux, Alain Monfort, and Eric Renault. Indirect Inference. *Journal of Applied Econometrics*, 8(S1):S85–S118, 1993.
- Lars Peter Hansen and Robert J. Hodrick. Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis. *Journal of Political Economy*, 88(5):829–853, 1980.
- Steven L. Heston. A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *Review of Financial Studies*, 6(2):327–343, 1993.
- Roni Israelov and Bryan Kelly. Forecasting the Distribution of Option Returns. *SSRN Working Paper*, 2017.
- Roni Israelov and Lars N. Nielsen. Covered Calls Uncovered. *Financial Analysts Journal*, 71(6):44–57, 2015.
- Bryan Kelly, Gerardo Manzo, and Diogo Palhares. Credit-Implied Volatility. *SSRN Working Paper*, 2019.
- Robert C. Merton. On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. *Journal of Finance*, 29(2):449–470, 1974.
- Robert C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1-2):125–144, 1976.
- Yoshio Nozawa. What Drives the Cross-Section of Credit Spreads?: A Variance Decomposition Approach. *Journal of Finance*, 72(5):2045–2072, 2017.
- Jun Pan. The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics*, 63(1):3–50, 2002.
- Pedro Santa-Clara and Shu Yan. Crashes, volatility, and the equity premium: Lessons from S&P 500 options. *Review of Economics and Statistics*, 92(2):435–451, 2010.
- Robert F. Stambaugh. Predictive regressions. *Journal of Financial Economics*, 54(3):375–421, 1999.