# Demand in the Option Market and the Pricing Kernel*

**Caio Almeida**[†] **Gustavo Freire**[‡]

This draft: January 31, 2023
First draft: January 31, 2022

**Abstract**

We show that net demand in the S&P 500 option market is fundamental to explain empirical puzzles related to the pricing kernel. When public investors (non-market makers) are exposed to variance risk by *net-selling* out-of-the-money (OTM) options, the pricing kernel is U-shaped, expected option returns are low and the variance risk premium is high. Conversely, when public investors are protected against variance risk by *net-buying* OTM options, the pricing kernel is decreasing in market returns, expected option returns are high and the variance risk premium is low. Our findings support equilibrium models with heterogeneous agents in which options are nonredundant.

*JEL classification:* G11, G12, G13

*Keywords:* Pricing Kernel, Option Returns, Option Demand, Market Makers, Risk Premium

*We would like to thank Yacine Ait-Sahalia, Hui Chen, Darrell Duffie, Mathieu Fournier, René Garcia, Felipe Iachan, Erik Kole, Raul Riva, David Schreindorfer, João Thereze, Michel van der Wel, Motohiro Yogo, conference participants at the 14th Annual SoFiE Conference, and seminar participants at the Erasmus School of Economics for useful comments and suggestions.

[†]E-mail: calmeida@princeton.edu, Department of Economics, Princeton University.

[‡]E-mail: freire@ese.eur.nl, Erasmus School of Economics, Erasmus University Rotterdam, and Tinbergen Institute.

# 1. Introduction

The pricing kernel, or stochastic discount factor (SDF), plays a central role in finance and economics as it describes how investors demand compensation for risk. Since Ait-Sahalia and Lo (2000), Jackwerth (2000) and Rosenberg and Engle (2002), index options have been used to recover the projection of the economy-wide pricing kernel onto market return states. Empirically, it is now a stylized fact that this projection is often not monotonically decreasing (Cuesdeanu and Jackwerth, 2018). In a representative investor framework where the market index proxies for aggregate wealth, the violation of monotonicity is puzzling as it contradicts risk averse behavior.

The pricing kernel puzzle is intimately related to puzzling patterns observed for index option returns. Expected returns of both put and call options are too low to be consistent with standard representative agent models (Coval and Shumway, 2001; Driessen and Maenhout, 2007; Bakshi, Madan and Panayotov, 2010; Bondarenko, 2014; Chaudhuri and Schroder, 2015; Baele et al., 2019). This means that, on average, the SDF is high for both large negative and positive market returns, following a U-shaped pattern. The low returns of index options are also manifested in a high premium for variance risk (Bakshi and Kapadia, 2003; Bollerslev, Tauchen and Zhou, 2009; Carr and Wu, 2009).

In this paper, we examine whether net demand in the S&P 500 option market helps explain the shape of the pricing kernel and puzzles related to it. Our motivation is simple. While it is counter-intuitive that large positive market returns represent a bad state (and, thus, high marginal utility) for a representative investor, it is intuitive that they represent a bad state for investors selling out-of-the-money (OTM) calls. Analogously, marginal utility for negative market return states would be especially high for investors selling OTM puts. If the representative agent assumption does not hold, the option-implied SDF reflects risk preferences of a marginal investor in the option market, such that net options positions can be informative about risk exposures and compensation for risk.

In fact, the representative investor framework is arguably troublesome to interpret the pricing kernel recovered from options. The representative agent's equilibrium demand for options must be zero as options are in zero net supply, such that the agent's marginal utility does not depend on the options. That is, this framework essentially assumes that options are redundant securities. This is inconsistent with the large trading volumes of index options and extant empirical evidence that options are nonredundant (Bakshi, Cao and Chen, 2000; Buraschi and Jackwerth, 2001; Almeida and Freire, 2022). If options are nonredundant, they should affect the functional form of the SDF, in a way that depends on how agents optimally trade the options (Vanden, 2004, 2006). Option demand arises with heterogeneous agents in the presence of frictions or sources of market incompleteness (Johnson, Liang and Liu, 2018).

Given our motivation, we investigate empirically the relation between net demand in the S&P 500 option market and the shape of the option-implied pricing kernel. There are essentially two types of agents in this market (Garleanu, Pedersen and Poteshman, 2009): public investors, or "end-users", who directly seek option exposure, and market makers, or "dealers", who provide liquidity by taking the other side of the end-user net demand. Dealers intermediate the market and must accommodate order imbalances from public investors, such that their options positions often deviate from the desired level (Muravyev, 2016). Consequently, market makers accumulate large inventories over time and cannot directly control their risk exposures (Gruenthaler, 2022).

We compute the net amount of new OTM calls and OTM puts that public investors buy each month, hencefort $PNBOC$ and $PNBOP$, respectively. This reflects the net amount that market makers sell in the same period. While on average public investors are net buyers of these options, they often switch to net sellers, especially during periods of market turmoil.[1] Chen, Joslin and Ni (2019) show that tight financial constraints often force market makers to aggressively hedge the left tail risk exposures of their inventories by buying OTM puts. We find a similar pattern for OTM calls, suggesting market makers are also interested in hedging market upside risk during crises.

To study the shape of the pricing kernel, we analyze the behavior of expected option returns across strike prices.[2] We calculate average call (put) returns for our entire sample and conditioned to periods where $PNBOC$ ($PNBOP$) is positive and negative. Our unconditional results are in line with the literature. Average call returns are mostly negative, providing evidence of a U-shaped pricing kernel. As for put options, average returns are highly negative, indicating that on average the SDF projection has a steep slope for negative market returns.

However, expected option returns are strikingly different depending on the option net demand. When public investors are exposed to market upside risk by net selling OTM

[1] More specifically, over our sample ranging from January 1996 to December 2020, $PNBOC$ ($PNBOP$) is negative for 52% (47%) of the 300 months.

[2] This approach is fully nonparametric and avoids the need to directly estimate the SDF projection, which is empirically challenging (see Coval and Shumway, 2001; Bakshi, Madan and Panayotov, 2010; Chaudhuri and Schroder, 2015; Baele et al., 2019).

calls, average call returns are more negative than for the whole sample, consistent with a U-shaped SDF that is steeper in the region of positive market returns. In contrast, average call returns are highly positive when public investors are net buyers of OTM calls, aligned with an overall monotonically decreasing projection of the pricing kernel.

Analogously, when public investors are more exposed to market downside risk by net selling OTM puts, average put returns are more negative than for the whole sample, i.e., the SDF projection is even steeper for negative market returns. For months where $PNBOP > 0$, average put returns are less negative and can even be highly positive if public investors are net buying large amounts of OTM puts. Our results are robust to controlling for several alternative determinants of the shape of the pricing kernel and expected option returns.

Given that option net demand strongly affects expected option returns, it should also be important to explain the variance risk premium.[3] We test that relation nonparametrically by analyzing expected delta-hedged option returns and straddle returns, which are informative about the premium for variance risk (Coval and Shumway, 2001; Bakshi and Kapadia, 2003). When public investors are exposed to variance risk by selling OTM options, the average returns of these strategies are highly negative, indicating a large positive variance risk premium.[4] In contrast, when public investors are protected against variance risk by buying OTM options, average returns are highly positive, consistent with a negative variance risk premium. We also run predictive time series regressions of the $VRP$ measure of Bollerslev, Tauchen and Zhou (2009) on net demand.[5] We find that both $PNBOC$ and $PNBOP$ significantly predict the $VRP$, with negative coefficients.

Overall, we document that net demand in the S&P 500 option market is fundamental to explain the shape of the pricing kernel, expected option returns and the variance risk premium. Our results suggest that public investors can be seen as the marginal investor in the option market, as their risk exposures align with the shape of the option-implied pricing kernel and the compensation for variance risk. This contrasts with the common view that market makers are always hedge-providers to end-users, demanding a high variance risk premium due to their risk exposure. In fact, the variance risk premium is

[3] For instance, Baele et al. (2019) show that the variance risk premium can be written as the expected return on a portfolio of OTM calls and OTM puts with different strike prices.

[4] As Bollerslev, Tauchen and Zhou (2009) and Baele et al. (2019), we define the variance risk premium as the difference between the risk-neutral and the physical expected variance of the market return.

[5] The $VRP$ is calculated as the difference between the squared VIX and the realized variance of S&P 500 returns over the previous month.

high when public investors provide hedge against variance risk to market makers. This usually happens in periods of distress where market makers are financially constrained and willing to pay high prices for options to reduce the risk of their inventories.[6]

More generally, our findings support equilibrium models with heterogeneous agents in which options are nonredundant. A recent example of such approach is by Farago, Khapko and Ornthanalai (2021), who show that a heterogeneous-agent model incorporating asymmetric preferences rationalizes the dynamics of the put option market. Our analysis underscores that rationalizing the market for call options is also fundamental to explain empirical puzzles in the option market. Relatedly, incorporating these features could also be useful to resolve inconsistencies between sources of risk premia implied by options and leading asset pricing models (Beason and Schreindorfer, 2022).

The remainder of the paper is organized as follows. After a brief review of the related literature, we discuss the theoretical framework in Section 2. In Section 3, we describe our data and empirical design. Section 4 contains the main empirical results, while Section 5 reports the robustness analysis. Section 6 concludes the paper.

## *1.1. Related literature*

Our paper is mainly related to four strands of the literature. The first strand is the literature on the pricing kernel puzzle, initiated by Ait-Sahalia and Lo (2000), Jackwerth (2000) and Rosenberg and Engle (2002), who were the first to estimate the projection of the SDF from option data and document its nonmonotonic shape. The puzzle has been empirically confirmed by several papers afterwards, with proposed explanations that range from heterogeneous beliefs (Ziegler, 2007; Bakshi, Madan and Panayotov, 2010), market incompleteness (Hens and Reichlin, 2013) and additional state variables (Chabi-Yo, Garcia and Renault, 2008; Brown and Jackwerth, 2012; Christoffersen, Heston and Jacobs, 2013) to behavioral models (Polkovnichenko and Zhao, 2013; Baele et al., 2019).[7] We show that net demand in the option market is fundamental to explain the shape of the pricing kernel. Departing from a representative investor framework, a non-

[6]This is consistent with evidence from Gruenthaler (2022) that, in times of elevated distress, option sell prices change more than buy prices, suggesting that market makers want to induce public sell orders. Johnson, Liang and Liu (2018), Chen, Joslin and Ni (2019) and Gruenthaler (2022) argue that option intermediaries' appetite to hold risks is normally high but decreases around crises, where the risks are transferred to public investors.

[7]See also Barone-Adesi, Engle and Mancini (2008), Beare and Schmidt (2016), Song and Xiu (2016), Grith, Hardle and Kratschmer (2017), Sichert (2022) and references in Cuesdeanu and Jackwerth (2018).

monotonic projection of the SDF onto market returns is not puzzling, but rather reflects the marginal investor's risk exposures which depend on options positions. More generally, any alternative explanation for the pricing kernel puzzle should be able to explain the stylized facts we document.

The second strand of the literature studies empirical patterns in index option returns, documenting low put returns, low call returns and relating expected option returns to the shape of the pricing kernel and the variance risk premium (Coval and Shumway, 2001; Bakshi and Kapadia, 2003; Jones, 2006; Driessen and Maenhout, 2007; Santa-Clara and Saretto, 2009; Bakshi, Madan and Panayotov, 2010; Bondarenko, 2014; Chaudhuri and Schroder, 2015; Baele et al., 2019). We document that expected option returns are not always low, but rather differ substantially once we condition to periods where public investors are net buyers or net sellers of OTM options.

The third strand of the literature investigates the role of demand in the option market in explaining asset prices and risk premium. Bollen and Whaley (2004) show that changes in implied volatility are correlated with signed option volume, while Garleanu, Pedersen and Poteshman (2009) model the effects of demand pressure on option expensiveness. Muravyev (2016) analyzes the implications of option order imbalances to equity option returns. More recently, Chen, Joslin and Ni (2019) show that supply shocks in deep OTM puts predict excess returns for a variety of assets. Fournier and Jacobs (2020) develop an option pricing model incorporating the intermediation of risks by market makers, while Jacobs, Mai and Pederzoli (2022) estimate latent demand and supply for index options. Gruenthaler (2022) studies how option intermediaries' risk affect risk premia. We show that the net demand for S&P 500 options is a key ingredient to explain the shape of the pricing kernel, expected option returns and the variance risk premium.

The fourth is a growing literature focusing on the pricing kernel of specialized institutions and financial intermediaries, rather than that of a representative household, to explain asset pricing behavior (Gabaix, Krishnamurthy and Vigneron, 2007; He and Krishnamurthy, 2013; Adrian, Etula and Muir, 2014; Brunnermeier and Sannikov, 2014; He, Kelly and Manela, 2017). We consider the Euler equation for a marginal investor in the option market instead of the one for a representative investor. Our results suggest that public investors in the S&P 500 option market, which are mostly sophisticated institutional investors (Chen, Joslin and Ni, 2019), can be seen as marginal investors.

# 2. Theoretical framework

The absence of arbitrage is equivalent to the existence of a strictly positive random variable $M_{t,T}$, the pricing kernel, such that the current price $P_t$ of any asset with future payoff $X_T$ satisfies:[8]

$$P_t = \mathbb{E}_t[M_{t,T} X_T]. \tag{1}$$

In traditional consumption-based asset pricing models (e.g., Lucas, 1978), the SDF equals the intertemporal marginal rate of substitution between consumption at time $t$ and $T$. More generally, however, the pricing kernel can be a function of multiple state variables that affect marginal utility. To avoid the issue of specifying these variables, a common approach is to consider the projection of the SDF onto market returns $R_{t,T}$:

$$m_{t,T}(R_{t,T}) = \mathbb{E}_t[M_{t,T} | R_{t,T}]. \tag{2}$$

The projected pricing kernel $m_{t,T}$ has the same pricing implications as the economy-wide SDF $M_{t,T}$ for assets with payoffs dependent on $R_{t,T}$.

The SDF induces a change of measure from the physical measure $\mathbb{P}$ to the risk-neutral measure $\mathbb{Q}$ under which one can take simple expectations to calculate the price of an asset. More specifically, the projected pricing kernel is the ratio of the conditional risk-neutral and physical distributions of market returns, discounted by the risk-free rate $R_f$:

$$m_{t,T}(R_{t,T}) = \frac{1}{R_f} \frac{\mathrm{d}\mathbb{Q}_t(R_{t,T})}{\mathrm{d}\mathbb{P}_t(R_{t,T})}. \tag{3}$$

The literature following Ait-Sahalia and Lo (2000), Jackwerth (2000) and Rosenberg and Engle (2002) calculates the empirical projection of the pricing kernel based on the equation above. The risk-neutral distribution is obtained from the cross-section of index option prices at date $t$ with time to maturity $T$, based on classic spanning results from Breeden and Lizenberger (1978). The physical distribution is estimated from the historical time series of market returns.[9]

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

[9]There is no consensus on how to estimate a conditional physical distribution, which is empirically challenging. Ait-Sahalia and Lo (2000) and Jackwerth (2000) use kernel densities to smooth the histogram of recent past market returns, while Rosenberg and Engle (2002), Barone-Adesi, Engle and Mancini (2008), Christoffersen, Heston and Jacobs (2013), Sichert (2022) and Almeida and Freire (2022) obtain conditional estimates taking into account time-varying volatility. Linn, Shive and Shumway (2017) and Barone-Adesi et al. (2020) combine option data with market returns to estimate the physical distribution.

## 2.1. *The pricing kernel puzzle*

To be able to economically interpret the projected pricing kernel, most of the literature assumes a representative investor framework where the market index is a perfect proxy for the aggregate wealth. Under these assumptions, the projection of the SDF onto market returns is equal to the economy-wide SDF, which in turn is proportional to the marginal utility of the representative investor. Empirically, it has been widely documented that this projection is often not a monotonically decreasing function of wealth as proxied by the market index (Cuesdeanu and Jackwerth, 2018). This would imply that the representative agent is locally risk-seeking, contradicting standard economic theory. Such violation of monotonicity characterizes the pricing kernel puzzle (Brown and Jackwerth, 2012).

The interpretation above implicitly assumes that options are redundant securities. Since options are in zero net supply, the representative agent's marginal utility does not depend on the options. Therefore, options are merely useful tools for the estimation of the pricing kernel, playing no role in its interpretation. This is at odds with the tremendous growth experienced by the index option market over the last few decades. In particular, the demand for S&P 500 options represents around 26% of total index shares in the recent sample, which confirms the relative importance of this market in the economy (Farago, Khapko and Ornthanalai, 2021).

There is also extensive empirical evidence that options are nonredundant. Bakshi, Cao and Chen (2000) document that call option prices can decrease even if the underlying asset price increases. Buraschi and Jackwerth (2001) show that option returns are not spanned by S&P 500 returns, which relates to the widely documented presence of additional priced risk factors such as stochastic volatility and jumps (Heston, 1993; Bates, 2000; Andersen, Fusari and Todorov, 2015). Ait-Sahalia, Wang and Yared (2001) find that the risk-neutral dynamics inferred from option prices is different from that extracted from S&P 500 returns. Almeida and Freire (2022) show that there is no unique well-behaved SDF estimated from S&P 500 returns that prices all index options in the cross-section.

If options are nonredundant, they should affect the economy-wide SDF, in a way that depends on how heterogeneous agents optimally trade them. As discussed by Johnson, Liang and Liu (2018), the main reason for option trading in equilibrium models is for transferring unspanned market risk (e.g., volatility or jump risk) across investors who can differ on preferences (Vanden, 2006; Bates, 2008; Bhamra and Uppal, 2009; Chen, Joslin and Ni, 2019), beliefs (Buraschi and Jiltsov, 2006), portfolio constraints (Vanden, 2004) or background risk (Franke, Stapleton and Subrahmanyam, 1998).

Therefore, option demand should be informative about the shape of the pricing kernel. In the next subsection, we propose a simple framework for interpreting the relation between options positions of the marginal investor in the option market and the SDF. This framework would be consistent with a general equilibrium model with heterogeneous agents and frictions. For instance, in the presence of nonnegative wealth constraints, Vanden (2004) shows that options are nonredundant and the economy SDF can be interpreted as the marginal utility of a pricing agent with smooth utility function trading in the market and options on the market.

## 2.2. *An alternative economic interpretation*

In this subsection, we depart from the representative investor framework and provide an alternative economic interpretation for the pricing kernel implied by index options. We start from the more general setup in which there is (at least) one marginal investor in the index and option markets. Let $S_t$ denote the market index and $O_{t,T,j}$ the prices of $j = 1, ..., N-1$ OTM index options with strike price $K_j$ expiring at time $T$.[10] Given the market returns $R_{t,T}$, the option returns are defined by:

$$f(R_{t,T}, K_j) = \begin{cases} \frac{(S_t R_{t,T} - K_j)^+}{O_{t,T,j}} & \text{if } K_j > S_t \text{ (OTM calls)} \\ \frac{(K_j - S_t R_{t,T})^+}{O_{t,T,j}} & \text{if } K_j < S_t \text{ (OTM puts).} \end{cases} \tag{4}$$

We stack the returns of the market and the options on an $N$-dimensional random vector $\mathbf{R}$ and denote by $\mathbf{R}^e \equiv \mathbf{R} - R_f \mathbf{1}_N$ the vector of excess returns with respect to the risk-free rate, where $\mathbf{1}_N$ is a conformable vector of ones.[11] In this setting, the marginal agent's SDF $m$ satisfies the Euler equation for the market and the options:

$$\mathbb{E}[m\mathbf{R}^e] = \mathbf{0}_N. \tag{5}$$

To shed light on the relation between option demand of the marginal investor and the SDF, we consider a standard optimal portfolio problem where the investor has a general concave utility function $U(W)$ defined over next period's wealth $W$. The investor

[10] This is without loss of generality in the sense that, from put-call parity, in-the-money (ITM) options contain redundant information. This is also consistent with the practice of using only OTM options to estimate the pricing kernel projection, as they are more liquid and actively traded than ITM options.

[11] For ease of notation, we drop time subscripts from now on.

distributes her initial wealth $W_0$ by investing $\lambda = (\lambda_1, ..., \lambda_N)$ in the market and the options and the remaining $W_0 - \sum_{j=1}^{N} \lambda_j$ in the risk-free asset. Terminal wealth is given by $W(\lambda) = W_0 R_f + \lambda \mathbf{R}^e$ and the investor maximizes expected utility:

$$\lambda^* = \underset{\lambda \in R^N}{max} \ \mathbb{E}\left[U(W(\lambda))\right]. \tag{6}$$

The marginal agent's first-order condition gives the Euler equation for the market and the options:

$$\mathbb{E}[U'(W_0 R_f + \lambda^* \mathbf{R}^e)\mathbf{R}^e] = \mathbf{0}_N, \tag{7}$$

such that the marginal utility of the investor as a function of her optimal future wealth serves as the SDF: $m(\mathbf{R}) = U'(W_0 R_f + \lambda^* \mathbf{R}^e)$.

Given the concavity of the utility function, the investor is risk-averse and $m^*(\mathbf{R})$ is monotonically decreasing in the optimal portfolio returns $\lambda^* \mathbf{R}^e$. However, this is not necessarily true for the projection of $m^*(\mathbf{R})$ onto market returns. In fact, if options are nonredundant, the endogenous optimal wealth $W(\lambda^*)$ will in general be different from the market index. In this context, the projection of the SDF onto market return states is not puzzling regardless of its shape. Rather, its shape simply reflects how the marginal investor trades in the options, as we discuss in the next subsection.

Our interpretation is also motivated by how the projection of the pricing kernel onto market returns is estimated in practice. A common way of estimating the numerator in (3) is to choose the risk-neutral probabilities that minimize a convex loss function while correctly pricing the market index and the options (Jackwerth and Rubinstein, 1996). In Appendix B, we show that this is equivalent, by duality, to solving precisely the optimal portfolio problem considered here. That is, using options to estimate the pricing kernel is equivalent to identifying the SDF of a marginal investor trading in the market index and the index options. This perspective further shows how troublesome it is to ignore the role of options in the shape of an SDF estimated from options.

## *2.3. Implications for the shape of the pricing kernel projection*

The marginal investor's pricing kernel $m^*(\mathbf{R})$, or, equivalently, her marginal utility, is high (low) for low (high) levels of endogenous wealth, i.e., for negative (positive) realizations of the optimal portfolio returns. The shape of the pricing kernel projection onto market returns will therefore reflect how market returns relate to the optimal portfolio returns. This, in turn, depends on the options positions in the portfolio. In the simplest case possible where portfolio weights for all options are zero, the optimal portfolio is perfectly correlated with the market index, so that the SDF is monotonically decreasing in market returns. However, in the more general case that options are nonredundant, the market index will differ from the endogenous wealth arising from the optimal portfolio.

We now derive necessary conditions on the investor's optimal portfolio to generate nonmonotonicity of the projection of the pricing kernel onto market returns. The next proposition focuses on the shape of the projected SDF in the upside of market returns:

**Proposition 1.** *Suppose the investor is long in the market index.*[12] *Then, if the projection of the pricing kernel onto market returns is U-shaped, the investor must be selling OTM call options. Otherwise, the projected pricing kernel is monotonically decreasing in the region of positive market returns.*

*Proof.* See Appendix A. □

The result above is fairly intuitive. In the region of positive net market returns, the only options that pay off are OTM calls, such that only their returns are informative about the shape of the SDF projection. In particular, positive market returns are associated with positive OTM call returns. Therefore, if the investor is long in all OTM call options, positive market return states represent positive realizations of the optimal portfolio and the projected SDF is monotonically decreasing. The projected SDF is U-shaped when the investor sells a sufficiently large amount of OTM calls that offsets the long position in the market. In this case, the positive call returns in the region of positive market returns will represent negative realizations of the optimal portfolio, i.e., states of higher marginal utility. That is, the pricing kernel is U-shaped when the investor is exposed to market upside risk by selling OTM calls.

In the next proposition, we consider the case that the pricing kernel is increasing in the downside region of market returns:[13]

**Proposition 2.** *Suppose the investor is long in the market index. Then, if the projection of the pricing kernel is monotonically increasing in the region of negative market returns,*

[12]This is the natural starting point where, without the options, the pricing kernel would be monotonically decreasing in the market returns. If the investor is short in the market index, the pricing kernel trivially violates the monotonically decreasing condition.

[13]This type of nonmonotonicity would be related to S-shaped pricing kernels that have also been found empirically in the literature (see Cuesdeanu and Jackwerth, 2018).

*the investor must be buying OTM put options. Otherwise, the projected pricing kernel is monotonically decreasing in this region.*

*Proof.* See Appendix A. □

The intuition is similar to the previous case. In the region of negative net market returns, only OTM puts pay off, such that only their returns are informative about the shape of the SDF projection. In particular, negative market returns are associated with positive OTM put returns. Therefore, if the investor is short in all OTM put options, negative market return states represent negative realizations of the optimal portfolio and the projected SDF is monotonically decreasing. The projected SDF is increasing in this region when the investor buys a sufficiently large amount of OTM puts that offsets the long position in the market. In this case, the positive put returns in the region of negative market returns will represent positive realizations of the optimal portfolio, i.e., states of lower marginal utility. That is, the investor profits from market downside risk.

The implications of the options positions of the marginal investor in the option market for the shape of the projected pricing kernel can also be cast in terms of implications for expected option returns and the variance risk premium. We briefly discuss these implications in the next subsection.

### *2.4. Implications for expected option returns and variance risk premium*

There is a close relation between the shape of the pricing kernel projected onto market returns and expected index option returns. Coval and Shumway (2001) show that the SDF is a monotonically decreasing function of market returns if and only if expected call (put) option returns are positive (negative) and increasing in the strike price. Instead, if the SDF is U-shaped, Bakshi, Madan and Panayotov (2010) demonstrate that expected call returns are decreasing in the strike and negative beyond a threshold. The steeper the slope of the pricing kernel in the region of positive (negative) market returns, the more negative are the expected call (put) returns.[14]

Given the relation above, Proposition 1 implies that if the marginal investor buys all OTM calls, expected call returns are positive and increasing in the strike. In contrast, if expected call returns are decreasing in the strike and negative beyond a threshold, the investor must be selling OTM calls. Similarly, Proposition 2 states that if the marginal

[14]For a more detailed explanation and a visual illustration, see Appendix C.

investor sells all OTM puts, expected put returns are negative and increase with the strike. If, instead, expected put returns are positive and decrease with the strike, the investor must be buying OTM puts.

Moreover, Baele et al. (2019) show that the variance risk premium can be written as a weighted average of expected returns on OTM calls and OTM puts, with negative weights. Therefore, our theory implies that if the marginal investor mostly sells OTM calls and puts, expected call and put returns are negative, and thus the variance risk premium should be positive. Conversely, if the investor mostly buys OTM calls and OTM puts, expected call and put returns are positive, and the variance risk premium should be negative. In other words, there is a negative relation between the net amount of OTM options that the investor buys and the variance risk premium. Intuitively, by selling OTM calls and OTM puts, the investor gets exposed to market upside risk and more exposed to downside risk, respectively. This increases the exposure to variance risk, such that the investor requires a higher premium to bear it.

In sum, our theoretical framework connects (nonmonotonic) patterns in the shape of the SDF projection to the options positions of the marginal investor in the option market. In the next section, we describe our approach to empirically investigate the relation between net demand in the S&P 500 option market and the pricing kernel. The idea is that, if options are nonredundant, net demand should be informative about investors' risk exposures and compensation for risk, and, therefore, about the shape of the pricing kernel projection onto market returns. The framework provided here guides our interpretation of the empirical patterns we document.

# 3. Data and empirical design

## *3.1. Net demand in the S&P 500 option market*

To compute net demand in the S&P 500 option market, we use the Chicago Board Options Exchange (CBOE) Open-Close database, which contains the daily non-market maker option volume differentiated by who initiates the trade (customer or firm) and the type of the transaction (open-buy, open-sell, close-buy or close-sell).[15] Customers, or public investors, include retail investors and institutional investors such as hedge funds,

[15] More specifically, the types of transaction are: open-buy, a buy to open a new long position; open-sell, a sell to open a new short position; close-buy, a buy to close an existing short position; and close-sell, a sell to close an existing long position.

while firms include securities broker-dealers who are not designated market makers. As broker-dealers are more similar to market makers, we follow common practice of merging both into one group that trades against the public investors. We also follow Lemmon and Ni (2014) and Chen, Joslin and Ni (2019) in focusing on open orders, as close orders are mechanically influenced by existing positions and do not solely reflect investors' perceptions of the future. Our sample ranges from January, 1996 to December, 2020.

We construct two measures of net buying-to-open volume of public investors: one for OTM calls and one for OTM puts. According to our theoretical framework, these are the options that are informative about the shape of the projected pricing kernel in the region of positive and negative net market returns, respectively. We define OTM calls (puts) as call (put) options with moneyness $K/S \geq 1.03$ ($K/S \leq 0.97$). We use volume data for options with time to maturity close to one month, which is the horizon considered by the literatures on the pricing kernel puzzle and option returns. More specifically, we focus on short-term OTM options with time to maturity between 15 and 60 calendar days. Considering these options, the public net buying-to-open of OTM calls ($PNBOC_t$) is defined as the difference between the public open-buy volume and public open-sell volume of OTM calls in month $t$. The definition of the public net buying-to-open of OTM puts ($PNBOP_t$) is analogous. The public net buy reflects the net amount that market makers sell in the same period.

Garleanu, Pedersen and Poteshman (2009) document that public investors are net buyers of index options on average. More recently, Chen, Joslin and Ni (2019) show that the public net demand of deep OTM put options varies substantially over time, being mostly negative between the 2008 financial crisis and the 2011 European debt crisis. They argue that tight financial constraints force market makers to aggressively hedge the left tail risk exposures of their inventories by becoming net buyers of deep OTM puts. The authors also provide evidence that institutional public investors, which are predominant in the S&P 500 option market compared to retail investors, are the ones satisfying the market maker demand in periods of market distress.

Figure 1 plots the time series of $PNBOC$ and $PNBOP$. As can be seen, both $PNBOC$ and $PNBOP$ vary considerably over time, being negative for roughly half of the sample.[16] The correlation of the measures is 37.85% and they display some similarities. This is especially true between 2016 and 2019, where public investors are consistently net

[16]In fact, the average $PNBOC$ ($PNBOP$) is close to zero at the scale of the plot: $0.02 \times 10^5$ ($0.04 \times 10^5$).

buyers of OTM calls and OTM puts. This changes during the COVID crisis, as $PNBOC$ becomes negative. Before this period, patterns are less comparable. From the beginning of the sample, $PNBOC$ oscillates from positive to negative, until it becomes negative for most of the period between 2007 and 2012. Interestingly, public investors become net sellers of OTM calls right before the global financial crisis. In contrast, $PNBOP$ is mostly positive until the peak of the financial crisis in October, 2008. After that, public investors persistently net sell OTM puts to market makers until 2016.

While explaining the time variation in the net demand of OTM calls and OTM puts is not the focus of our paper, a similar argument as in Chen, Joslin and Ni (2019) may apply. In normal times, market makers primarily net sell OTM options, accumulating both left and right tail risk exposures in their inventories. Figure 1 shows that public investors become net sellers mostly when the market is in distress. These are periods where market makers face tight intermediary constraints, such as loss of capital and higher capital requirements. Therefore, to hedge the exposures of their inventories to variance risk, intermediaries switch to net buyers of OTM options. In Section 5.1, we provide additional insights by relating option net demand to several risk variables, indicators of business conditions and measures of funding constraints.

### *3.2. S&P 500 option returns*

To investigate the implications of public option net demand for the shape of the pricing kernel, we rely on the relation between the latter and expected option returns. This approach is fully nonparametric and avoids the need to explicitly calculate the option-implied SDF, which depends on the challenging estimation of the conditional physical distribution and can be misspecified. While this relation holds for expected returns conditional on time $t$, Coval and Shumway (2001), Bakshi, Madan and Panayotov (2010) and Chaudhuri and Schroder (2015) condition down by taking unconditional expectations of option returns.[17] We proceed in a similar way, but taking expectations on different subsamples of option returns. More specifically, motivated by our theory, we compute average call and put returns conditioned to months where $PNBOC$ and $PNBOP$, respectively, are: very large (above one standard deviation), positive, negative and largely negative (below minus one standard deviation). In this sense, our results speak to the overall or "average" shape of the pricing kernel under these different conditioning sets.

[17]For more details, see Appendix C.

To calculate the option returns, we obtain end-of-day midquotes of S&P 500 options from OptionMetrics from January, 1996 to December, 2020. We follow Baele et al. (2019) in using interpolated prices of call (put) options with fixed deltas ranging from 0.2 to 0.8 (-0.8 to -0.2) and fixed time to maturity equal to 30 calendar days. OptionMetrics generates these prices by interpolating the implied volatility surface on each day. This has the advantage that for each day in the sample we have option observations with exactly the same delta and time to maturity. Thus, for each day and option delta, we construct the hold-until-maturity option returns by taking the ratio of the realized option payoff and the option price.[18] That is, we have daily overlapping option returns, which guarantees a large number of observations when we compute average returns under different conditioning sets.[19] We have tested using the traded option quotes instead of the interpolated ones to calculate option returns, and results are very similar.[20]

We also investigate the implications of option net demand for the variance risk premium, which is closely related to the shape of the pricing kernel. More specifically, we compute expected returns of option strategies that are informative about the variance risk premium. This approach, again, is fully nonparametric. Bakshi and Kapadia (2003) show that negative expected delta-hedged option returns correspond to a positive variance risk premium. Using the data described above, we calculate, for each day for the at-the-money (ATM) call and put options, the hold-until-maturity return of buying the option and shorting delta shares of the S&P 500 index, such that the net investment earns the risk-free rate.[21] Then, we analyze average returns conditioned to periods where public investors are net buyers of OTM options (i.e., protected against variance risk) or net sellers of OTM options (i.e., exposed to variance risk).

Relatedly, negative expected straddle returns signal a positive variance risk premium (Coval and Shumway, 2001; Driessen and Maenhout, 2007). For each day, we calculate the hold-until-maturity return of a long ATM straddle position, which consists of a simultaneous purchase of an ATM call and ATM put option. This strategy profits from increases in market volatility, such that a negative expected return signals a premium

[18]To calculate the realized option payoff, we use daily data on the S&P 500 index from Bloomberg.

[19]For instance, if we condition to months where $PNBOC > 0$, we have 21 observations for each month $t$ where $PNBOC_t > 0$.

[20]In this case, options have varying moneyness and time to maturity across different days. Therefore, for each day and for a given fixed target moneyness $y^*$, we select the option with time to maturity closest to 30 calendar days and moneyness closest to $y^*$, and use this observation to calculate the option return.

[21]More precisely, following Bakshi and Kapadia (2003), the delta-hedged return of the ATM call option with price $O_{t,T}$ is defined as $\pi/O_{t,T}$, where $\pi = max(S_T - K, 0) - O_{t,T} - \Delta_t(S_T - S_t) - r_t^f(O_{t,T} - \Delta_t S_t) \times 30$ and $r_t^f$ is the daily 3-month T-bill. The ATM call (put) is the one with $\Delta_t$ equal to 0.5 (-0.5).

for protection against variance risk. We also consider a "crash-neutral" ATM straddle, consisting of a long position in the ATM straddle and a short position in the -0.2 delta OTM put option. The return of this strategy is limited during a market crash, such that its expected return reflects compensation only for variance risk. That is, it isolates the premium for variance risk from that for crash risk. Again, we compute conditional average straddle returns based on whether public investors are net buying or net selling different quantities of OTM options.

### *3.3. Additional variables*

We use a number of additional variables in our empirical analysis. To conduct predictive time series regressions for the variance risk premium, we adopt a popular semi-parametric measure for the premium. We follow Bollerslev, Tauchen and Zhou (2009) in calculating the $VRP_t$ as the difference between the squared VIX index at month $t$ and the (annualized) realized S&P 500 return variance over $t-1$ and $t$. We obtain the VIX from the CBOE and the high-frequency five-minute S&P 500 returns, which we use to estimate realized variance (as the sum of the squared log returns), from Tick Data.

We also consider a set of variables which we use as controls in our tests and relate to $PNBOC_t$ and $PNBOP_t$ to provide further insights into net demand. First, we include risk measures. We consider the realized volatility ($RV_t$) as the square root of the realized variance over month $t$ and the left tail volatility ($LTV_t$) of Bollerslev, Todorov and Xu (2015), which we obtain from the website of Torben Andersen and Viktor Todorov.[22] Second, we include macroeconomic variables reflecting expected business conditions. Using updated data from Welch and Goyal (2008), we calculate: the term spread ($TMS_t$), as the difference between the long term (10-year) yield on government bonds and the 3-month Treasury bill; the credit spread ($CRS_t$), as the difference between the yield on BAA-rated corporate bonds and the long term yield on government bonds; and the default spread ($DFS_t$), as the difference between BAA and AAA-rated corporate bond yields.[23] We further obtain the Aruoba, Diebold and Scotti (2009) ($ADS_t$) business conditions index from the Federal Reserve Bank of Philadelphia. We also construct a dummy variable that equals one during NBER recession periods and zero otherwise ($D_t^{NBER}$).

Third, we include measures of funding constraints. We obtain the TED spread

[22]https://tailindex.com/index.html. Since $LTV_t$ is provided at a daily frequency, we calculate monthly figures by taking the average within each month. $LTV_t$ is only available until December, 2019.

[23]The data is available at Amit Goyal's website: http://www.hec.unil.ch/agoyal/.

($TED_t$), which is the difference between the 3-month LIBOR and the 3-month Treasury bill, from the St. Louis Federal Reserve Economic Data (FRED) database. We also construct the measure of financial intermediary constraints of Chen, Joslin and Ni (2019) ($CJN_t$). Their measure is defined as the monthly public net buying-to-open of deep OTM puts (with $K/S \leq 0.85$) times an indicator function that equals one for months where supply shocks are the dominant driver of price-quantity relations. Within each month $t$, a regression of daily $VRP$ on the daily public net buying-to-open of deep OTM puts is conducted to estimate the slope coefficient $b_{VPR,t}$.[24] If $b_{VPR,t} < 0$, there is evidence of a negative price-quantity relation, such that month $t$ is interpreted as being dominated by supply shocks, rather than demand shocks.

# 4. Main empirical results

## *4.1. Option net demand and the shape of the pricing kernel*

To study the shape of the pricing kernel, we analyze expected option returns across strike prices. We calculate average returns over the whole sample and conditioned to different months based on option net demand. Our main findings are summarized in Figure 2. Focusing first on the unconditional average returns, our results are in line with the literature. Average call returns are decreasing in the strike and negative beyond a threshold, providing evidence of a U-shaped pricing kernel. As for put options, average returns are increasing in the strike and highly negative, indicating that on average the SDF projection has a decreasing steep shape for negative market returns.

However, expected option returns are strikingly different depending on the public option net demand. The left panel of Figure 2 shows that when public investors are exposed to upside risk by net selling OTM calls, average call returns are more negative than for the whole sample, consistent with a U-shaped SDF that is steeper in the region of positive market returns. In contrast, average call returns are positive and increasing in the strike when public investors are net buyers of OTM calls, aligned with an overall monotonically decreasing projection of the pricing kernel. These differences are even more extreme when public investors are net buying or net selling large amounts of OTM calls.

The right panel of Figure 2 reports the patterns for put options. When public investors

[24] Daily $VRP$ is defined as the squared VIX for day $i$ minus the (annualized) realized S&P 500 return variance over the previous 30 calendar days.

are more exposed to downside risk by net selling OTM puts, average put returns are even more negative, i.e., the SDF projection is even steeper. This is exacerbated when $PNBOP < -1\text{std}$, where average put returns get extremely negative. On the other hand, for months where $PNBOP > 0$, average put returns are less negative as exposure to downside risk decreases. However, when public investors are protected against downside risk by net buying large amounts of OTM puts, average returns of put options are highly positive and decrease with the strike, suggesting an increasing SDF projection in the region of negative market returns.

The empirical results above show that the shape of the pricing kernel strongly depends on option net demand. This is consistent with our economic interpretation of the pricing kernel implied by index options. In particular, the compensation for risk reflected in the pricing kernel aligns with the risk exposures of public investors, suggesting that they can be seen as the marginal investor in the option market. Therefore, we can use our theoretical framework to interpret the relation between the options positions of public investors and the pricing kernel.

When public investors are net buyers of OTM calls, positive market returns represent "good" states of the world for them, such that the SDF is monotonically decreasing in this region. In contrast, when public investors net sell OTM calls, they are exposed to market upside risk and have higher marginal utility in positive market return states, leading to a U-shaped SDF. Moreover, when public investors net buy (sell) large amounts of OTM calls, exposure to upside risk decreases (increases), such that they require less (more) compensation for that risk with a steeper decreasing (increasing) SDF projection.

Analogously, when public investors are net sellers of OTM puts, they are more exposed to market downside risk, demanding more compensation for that risk with a steeper decreasing SDF projection. This is exacerbated when they net sell large amounts of OTM puts, as this increases even more the exposure to negative market return states. In contrast, exposure to downside risk is smaller when public investors net buy OTM puts.[25] However, when they net buy large amounts of OTM puts, negative realizations of the market represent positive realizations of the investors' portfolios and states of lower marginal utility. That is, the pricing kernel is monotonically increasing in the region of negative market returns.

[25]If the long positions in OTM puts are not enough to offset the long position in the market index, negative market returns are still painful to the investors, such that the SDF projection is still monotonically decreasing in this region, albeit with a flatter slope.

In sum, we document that the magnitude of average option returns and how they vary with the strike depend on the quantities of OTM options that public investors are net buying or net selling in the S&P 500 option market. Our findings indicate that the SDF projection in the region of positive (negative) market returns can be monotonically decreasing or increasing depending on the trading quantities of OTM calls (OTM puts). In particular, the compensation for risk reflected in the pricing kernel is in agreement with the risk exposures of public investors. In Appendix D, we report the results in further details and provide formal tests of the monotonicity of expected option returns across strikes. Importantly, we show that all the findings above are statistically significant.

### *4.2. Option net demand and the variance risk premium*

To study the relation between option net demand and the variance risk premium, we analyze expected delta-hedged option returns. Bakshi and Kapadia (2003) show that a negative (positive) expected return of this option strategy corresponds to a positive (negative) variance risk premium. We follow them in focusing on ATM options, which are the most sensitive to volatility and thus most informative about the variance risk premium. We calculate average returns over the whole sample and conditioned to periods where public net demand is above or below different thresholds for both OTM calls and OTM puts at the same time. The motivation is to investigate the compensation for variance risk when public investors are exposed to (protected against) this risk by net selling (buying) OTM options.

Figure 3 reports the results. Consistent with Bakshi and Kapadia (2003), the unconditional average delta-hedged returns of both calls and puts are negative ($-9.39\%$ and $-15.72\%$, respectively), indicating the presence of a positive variance risk premium in the whole sample. However, average returns change substantially once we condition to different quantities of net demand for OTM options. When public investors are net selling OTM options, average delta-hedged returns are even more negative, implying a larger variance risk premium. The opposite happens when these investors are net buying OTM options, where average returns actually become positive and large for higher amounts of net demand, reaching up to 49% for calls and 39% for puts. That is, the variance risk premium gets highly negative. These patterns are not only economically, but also statistically significant.

We also analyze the expected returns of ATM straddles. These option strategies essentially represent long positions in volatility and thus contain information about the variance risk premium (Coval and Shumway, 2001; Driessen and Maenhout, 2007). Our approach is analogous to that for the delta-hedged option returns. Figure 4 depicts the results. Unconditionally, both the ATM straddle and crash-neutral straddle earn negative average returns ($-12.56\%$ and $-3.37\%$, respectively), which is aligned with the previous literature. This means that there is a positive variance risk premium as a hedge against variance risk loses money on average.

However, when public investors are exposed to variance risk by net selling OTM options, average returns are even more negative and often twice as large in absolute value, signaling a positive and large variance risk premium. In contrast, when these investors are protected against variance risk by net buying OTM options, average returns become positive, suggesting a negative premium for this risk. The higher the quantities of OTM options that public investors net buy, the higher are the average returns of the straddle and crash-neutral straddle, reaching up to 43.81% and 31.14%, respectively. Again, these results are statistically significant. The evidence from the crash-neutral straddle further shows that our findings are not driven by compensation for crash risk.

The compensation for variance risk reflects the risk exposures of public investors, reinforcing the idea that they represent marginal investors in the option market. The patterns we find are, again, consistent with our theory. When public investors net sell OTM calls and OTM puts, they get exposed to market upside risk and more exposed to downside risk, respectively. That is, exposure to market variance risk increases and they require more compensation to bear it, such that the variance risk premium is higher. Conversely, when public investors net buy OTM calls and OTM puts, they profit from positive net market returns and are more protected against negative net market returns, respectively. This means that exposure to variance risk decreases, as does the premium to bear it. In fact, the premium gets even negative.

In sum, we provide new evidence that the variance risk premium is not generally positive and large, but depends on the trading quantities of public investors. The variance risk premium is positive (negative) when public investors are exposed to (protected against) variance risk by net selling (buying) OTM options. The compensation for variance risk is aligned with the risk exposures of public investors. In Section 5.3, we provide further evidence of the negative relation between option net demand and the variance risk premium using the semi-parametric $VRP$ measure.

### 4.3. Discussion

The empirical results in this section show that the shape of the pricing kernel, expected index option returns and the variance risk premium strongly depend on option net demand. This is in agreement with our economic interpretation of the pricing kernel implied by options. In particular, we find that the compensation for risk reflected in the pricing kernel aligns with the risk exposures of public investors instead of those of market makers, which is consistent with the structure of the option market. Public investors are the end-users who directly seek option exposure and trade in options to optimize their portfolios. In contrast, market makers intermediate the market and cannot directly control their options positions, as they must accommodate order imbalances from end-users (Muravyev, 2016). In fact, market makers' main source of income from option trading is the revenue from the bid-ask spread, which is their compensation for providing liquidity (Gruenthaler, 2022).

The risk exposures of public investors are also aligned with the observed variance risk premium. This is at odds with the common view that market makers are always providing hedge (against downside risk or variance risk) to end-users, demanding as such a high premium for their exposure to variance risk. In fact, Figure 1 shows that there are periods where public investors are actually net selling OTM options and thus providing hedge to market makers. We document that it is precisely during those periods that the variance risk premium is high. These are usually periods of market distress where market makers are constrained in their role as intermediaries and willing to pay high prices for OTM options to reduce the variance risk of their inventories.[26] This is consistent with evidence that option intermediaries' appetite to hold risks decreases around crises (Johnson, Liang and Liu, 2018; Chen, Joslin and Ni, 2019; Gruenthaler, 2022). In contrast, when public investors are protected against variance risk by net buying OTM options, the variance risk premium is low or even negative. These are usually periods where market makers' risk appetite is high and they can appropriately intermediate the market.

More generally, the stylized facts we document would be difficult to reconcile with representative agent models in which options play no role. Our findings support equilibrium models with heterogeneous agents in which options are nonredundant. In particular, while some attention has been given in the literature to model the market for OTM puts

[26] This is aligned with evidence from Gruenthaler (2022) showing that in times of elevated distress market makers try to induce public sell orders as sell prices change more than buy prices.

(Bates, 2008; Chen, Joslin and Ni, 2019; Farago, Khapko and Ornthanalai, 2021), our evidence highlights that taking OTM call options into account is also fundamental to explain empirical puzzles in the S&P 500 option market. Modeling option trading among heterogeneous agents may also help reproduce sources of risk premia implied by options that are inconsistent with leading asset pricing models (Beason and Schreindorfer, 2022).

# 5. Option net demand and alternative determinants

In this section, we regress $PNBOC_t$ and $PNBOP_t$ on a number of risk variables, business conditions indicators and funding constraints measures in order to provide additional insights into the time series of option net demand. We also show that the effects of option net demand on the shape of the pricing kernel and the variance risk premium are robust to and remain relevant after including several controls in the analysis.

## *5.1. Relation to risk, business conditions and funding constraints*

Table 1 reports the results from univariate contemporaneous regressions of $PNBOC_t$ and $PNBOP_t$ on different variables as indicated by the columns. Focusing on the first column, it can be seen that there is a highly significant negative relation between $PNBOC_t$ and the dummy variable for NBER recessions. This indicates that public investors tend to be net sellers of OTM calls during recessions. In contrast, there is a positive relation between $PNBOP_t$ and $D_t^{NBER}$. This is consistent with Figure 1 that shows that during recessions periods public investors were mostly net buying OTM puts. In fact, they switch to net sellers right after the global financial crisis and COVID crisis.

The second column contains the results for realized volatility. $PNBOC_t$ is strongly negatively related to $RV_t$, suggesting that during periods of high volatility public investors are net sellers of OTM calls. As for $PNBOP_t$, it is strongly positively related to $RV_t$, which resonates the previous result for $D_t^{NBER}$ as volatility is usually high during recessions. We also consider the $LTV_t$, which captures the expected volatility that stems from large negative price jumps, being thus related to jump tail risk. A strong negative relation between $PNBOC_t$ and $LTV_t$ suggests that when jump tail risk is high public investors tend to net sell OTM calls. The relation between $PNBOP_t$ and $LTV_t$ is insignificant.

We next focus on indicators of expected business conditions. The $TMS_t$ measures the slope of the yield curve and is often regarded as a pro cyclical variable: steepening of the slope signals expansion of the economy, while flattening indicates contraction. Table 1 shows that both $PNBOC_t$ and $PNBOP_t$ display a strong negative relation with $TMS_t$. This suggests that prior to expansions as predicted by $TMS_t$ public investors are mostly net selling OTM options. The $CRS_t$, on the other hand, is counter cyclical as it widens during weak business conditions to compensate investors for an increased probability of default. Only $PNBOC_t$ is significantly related to $CRS_t$, with a negative coefficient. That is, public investors tend to net sell OTM calls during periods of weak business conditions where $CRS_t$ is high. As for the $DFS_t$, it has a similar interpretation to $CRS_t$, but capturing credit riskiness within the corporate sector. Again, only $PNBOC_t$ is significantly related to $DFS_t$, where the coefficient is negative. The seventh column contains results for the $ADS_t$ index, which tracks real business conditions based on a number of underlying economic indicators. Positive (negative) values indicate better (worse) than average conditions. As can be seen, $ADS_t$ is only significantly related to $PNBOP_t$. The observed negative coefficient suggests that when business conditions are weak public investors tend to net buy OTM puts. This, again, resonates the fact that during NBER recessions $PNBOP_t$ was mostly positive.

Finally, we relate option net demand to measures of funding constraints. We focus first on the measure of Chen, Joslin and Ni (2019). A low $CJN_t$ is a sign of tight intermediary constraints, as it represents a low net demand for deep OTM puts due to supply shocks. Both $PNBOC_t$ and $PNBOP_t$ are strongly positively related to $CJN_t$, indicating that when option intermediaries are more constrained, public investors tend to be net sellers of OTM options. This is aligned with the idea that when market makers are financially constrained they need to hedge their tail risk exposures and thus become net buyers of OTM options. As for the $TED_t$, which measures the credit risk of banks, it is significantly negatively related to $PNBOC_t$. This indicates that when the default risk on interbank loans is high public investors are mostly net selling OTM calls. In contrast, there is a strong positive relation between $PNBOP_t$ and $TED_t$.

In sum, $PNBOC_t$ and $PNBOP_t$ are significantly related to a number of risk variables, macroeconomic indicators and funding constraints measures. These relations are interpretable and provide further insights into the dynamics of option net demand. In particular, while $PNBOC_t$ and $PNBOP_t$ share similarities, they also differ in a number of ways. Given such relations, one potential concern could be that option net demand is simply a proxy for alternative determinants of the shape of the pricing kernel and the variance risk premium. We address this concern in the next subsections, where we show that including these variables as controls in the analysis does not affect the explanatory power of option net demand for the empirical puzzles in the option market.

## 5.2. *Option net demand, controls and the shape of the pricing kernel*

In this subsection, we investigate the effects of option net demand on the shape of the pricing kernel after controlling for a number of alternative determinants. More specifically, we analyze how expected index option returns vary with the strike conditional on different levels of several control variables and net demand for OTM options.

### 5.2.1. *NBER Recessions*

We first investigate how the effects of option net demand on the shape of the pricing kernel change if we exclude the months indicated as recessions by the NBER. That is, we first condition on months where there is no recession, and then on different quantities of net demand. We analyze average option returns calculated under these different conditioning sets. Figure 5 reports the results. Excluding recessions, average call returns are slightly higher than for the whole sample, but still provide evidence of a U-shaped pricing kernel as average returns are decreasing in the strike and eventually negative. As for puts, average returns outside of recessions are more negative than for the whole sample, indicating a steeper monotonically decreasing pricing kernel in the region of negative market returns.

The left panel of Figure 5 shows that the effects of $PNBOC$ conditional on no recessions are essentially the same as for the whole sample. When public investors are exposed to upside risk by net selling OTM calls, average call returns are more negative, consistent with a steeper U-shaped SDF. In contrast, average call returns are highly positive and increasing in the strike when these investors are net buying OTM calls, indicating a monotonically decreasing projection of the pricing kernel on average. That is, the effects of $PNBOC$ do not depend on NBER recessions.

The right panel of Figure 5 plots the patterns for put options. When public investors are more exposed to downside risk by net selling OTM puts, results are the same as for the whole sample: average put returns are even more negative, providing evidence of a steeper monotonically decreasing SDF projection. Results are also similar when $PNBOP > 0$, where average returns are less negative as exposure to downside risk decreases. However, when $PNBOP > 1$std, it is no longer the case that average put returns are positive and decreasing with the strike. This means that the SDF is increasing in negative market returns only when public investors net buy large quantities of OTM puts during recessions. This result is not necessarily inconsistent with our theory. Outside of recessions, large trading quantities of OTM puts may not be enough to offset the long position in the market, such that the marginal investor is still exposed to downside risk and the pricing kernel is monotonically decreasing.

In sum, all our previous empirical results for the shape of the pricing kernel hold outside of recessions, with the exception of the nonmonotonicity of the pricing kernel in the region of negative market returns when $PNBOP > 1$std. This finding is informative as it tells us that such nonmonotonicity is concentrated in recession periods. That is, public investors only get effectively protected against downside risk when they are net buying large quantities of OTM puts during recessions.

### *5.2.2. Volatility*

We now investigate whether the effects of option net demand depend on volatility, which is arguably the most relevant variable to control for based on the literature. First, there is empirical evidence that the shape of the pricing kernel depends on volatility. Barone-Adesi et al. (2020) and Sichert (2022) show that standard empirical estimates of the pricing kernel implied by index options are U-shaped in periods of high volatility. On the other hand, Schreindorfer and Sichert (2022) document that the pricing kernel is steeper in the region of negative market returns during periods of low volatility.

Second, volatility has been considered as an additional state variable to explain the pricing kernel puzzle. Brown and Jackwerth (2012) posit that the pricing kernel is monotonically decreasing in states of low and high volatility but with a different slope, such that taking the expectation over volatility would yield a pricing kernel that is a nonmonotonic function of market returns. Christoffersen, Heston and Jacobs (2013) generalize the Heston and Nandi (2000) GARCH model by assuming that the pricing kernel depends not only on returns but also on volatility. Their parametric specification implies that the pricing kernel is a quadratic function of market returns, which by construction is U-shaped whenever the variance risk premium is positive.

We start by analyzing nonparametrically the relation between volatility and the shape of the pricing kernel. We compute average option returns for the whole sample and conditioned to months of very low, low, high and very high realized volatility. Figure 6 contains the results. Focusing first on the left panel, average call returns are more negative and decrease with the strike at a faster rate during periods of high volatility, providing stronger evidence of a U-shaped SDF compared to the whole sample. This confirms the semi-parametric evidence of Barone-Adesi et al. (2020) and Sichert (2022). During periods of low volatility, average call returns are higher and mostly positive but still decrease with the strike, suggesting a U-shaped SDF with a flatter slope on average. In contrast, average returns increase with the strike when volatility is very low, indicating a pricing kernel that is mostly monotonically decreasing for positive market returns.

The evidence above is inconsistent with the explanation of Brown and Jackwerth (2012). Their theory would imply that, once we condition to low or high volatility, the pricing kernel is monotonically decreasing but with different slopes, i.e., average call returns are positive and increase with the strike but at different rates. Instead, there is only evidence for a monotonically decreasing pricing kernel when volatility is very low. The GARCH model of Christoffersen, Heston and Jacobs (2013) rationalizes the U-shape with a fixed positive parameter for the variance risk premium. However, for the same fixed parameters, the model does not account for the fact that the pricing kernel can be monotonically decreasing (e.g., when $RV$ is very low).[27]

The right panel of Figure 6 reports the patterns for put options. When volatility is low, average put returns are more negative than for the whole sample, suggesting a monotonically decreasing pricing kernel that is steeper for negative market returns. In contrast, when volatility is high, average returns are less negative, indicating a flatter monotonically decreasing SDF projection. This confirms the semi-parametric evidence of Schreindorfer and Sichert (2022) that negative market returns are more painful to investors during periods of low volatility.

We next compute average option returns conditioned to option net demand after conditioning on high or low volatility states. The results appear in Figure 7. The effects of option net demand mostly survive after controlling for volatility. When volatility is

[27]It is also worth noting that the explanation of Christoffersen, Heston and Jacobs (2013) for the pricing kernel puzzle is conceptually very different from ours. They propose a parametric model for the pricing kernel that has a U-shape if the variance risk premium in the economy is positive. That is, a premium for variance explains the U-shape. In our case, we link observable measures of option net demand to the shape of the pricing kernel, which in turn determines expected option returns and the variance risk premium.

high, there is less evidence of a U-shaped pricing kernel when public investors are net buyers of OTM calls (average call returns are flatter for $PNBOC > 0$ and increasing in the strike for $PNBO >$ 1std). The results for puts during high volatility periods are essentially the same as those without controlling for $RV$: average put returns are more negative when public investors are net sellers of OTM puts, and are highly positive and decreasing with the strike when they are net buying large quantities of OTM puts.

The lower panels of Figure 7 display the patterns for low volatility periods. Results for calls are very similar to those without controlling for $RV$. When public investors are net selling OTM calls, average call returns are negative and strongly decrease with the strike, providing evidence of a steep U-shaped SDF. In contrast, when they are net buyers of OTM calls, average returns are positive and increasing in the strike, suggesting a monotonically decreasing SDF shape for positive market returns. As for the lower right panel, when volatility is low average put returns are highly negative and do not change much with $PNBOP$. This would be consistent within our framework with the long position in the market being dominant over put positions during low volatility periods.

In sum, we show that volatility is an important determinant of the shape of the pricing kernel, which complements previous empirical evidence from the literature. Even so, most of the effects of option net demand remain after controlling for $RV$. Overall, it is clear that option net demand provides unique information about the shape of the pricing kernel that is not captured by volatility and is consistent with our theoretical framework.

#### *5.2.3. Demand and supply shocks*

The distinction between demand and supply shocks plays an important role in the option demand literature. Garleanu, Pedersen and Poteshman (2009) show that demand shocks can generate a positive relation between option net demand and option expensiveness. Chen, Joslin and Ni (2019) show that such relation is negative during periods dominated by supply shocks. To the extent that option expensiveness can be related to the shape of the pricing kernel and expected option returns, it is pertinent to control for demand and supply shocks in our analysis. More specifically, our goal is to investigate whether the explanatory power of option net demand for the shape of the pricing kernel comes from capturing public investors' risk exposures or from being an indirect proxy for demand and supply shocks.

To that end, we follow the identification strategy of Chen, Joslin and Ni (2019) to determine if a given month is dominated by demand or supply shocks. Within each month $t$, we regress daily $VRP$ on the daily public net buying-to-open volume of deep OTM puts to estimate the slope coefficient $b_{VRP,t}$.[28] Deep OTM puts are used as they are more informative about intermediary constraints. If $b_{VRP,t} > 0$ ($b_{VRP,t} < 0$), there is evidence of a positive (negative) price-quantity relation, such that month $t$ is considered to be dominated by demand (supply) shocks. The higher the coefficient (in absolute value), the stronger the price-quantity relation.

Figure 8 plots average option returns across strikes conditioned to months dominated by demand shocks and months dominated by supply shocks. The left panel shows that average call returns are more negative during periods where supply shocks dominate, indicating a steeper U-shaped pricing kernel. When $b_{VRP} > 0$, average returns are positive but still decrease with the strike, suggesting a flatter U-shaped SDF. When the price-quantity relation is strongly positive, there is evidence of a monotonically decreasing pricing kernel. As for the right panel, it can be seen that average put returns do not change much with $b_{VRP}$. The exception is when there is a strongly positive price-quantity relation, where there is some evidence of nonmonotonicity of the SDF projection as average put returns decrease with the strike.

Figure 9 contains the results for option net demand after conditioning on demand or supply shocks. As can be seen from the left panels, regardless of demand or supply shocks, the pricing kernel is a steeper U-shaped function of market returns when public investors are exposed to upside risk by net selling OTM calls, while there is evidence for a monotonically decreasing pricing kernel when they are net buyers of OTM calls. As for the right panels, average put returns are more (less) negative when public investors are net sellers (net buyers) of OTM puts, both when $b_{VRP} > 0$ and $b_{VRP} < 0$. The main difference is that an increasing SDF projection in the region of negative market returns only appears when public investors are net buying large quantities of OTM puts during months dominated by demand shocks. This suggests that public investors only get protected against downside risk when market makers can appropriately provide liquidity, as they are mostly unconstrained when demand shocks prevail.

In sum, the relation between option net demand for OTM options and expected option returns is mostly unaffected by demand and supply shocks. This provides further evidence that the shape of the pricing kernel reflects compensation for the risk exposures of public

[28]Variables are standardized to have unit variance such that $b_{VRP,t}$ is comparable over time.

investors' new options positions. We also find that the nonmonotonicity of the pricing kernel in the region of negative market returns when $PNBOP > 1\text{std}$ is concentrated in periods dominated by demand shocks.

#### *5.2.4. Business conditions*

To further investigate the robustness of our findings to different macroeconomic conditions, we consider as a control the $ADS$ real business conditions index. We also have results for the remaining risk, macroeconomic and funding constraints measures in Section 5.1, which we omit for brevity. Importantly, the relation between option net demand and the shape of the pricing kernel is robust to controlling for all these other variables.

Figure 10 plots average option returns conditional on high or low values of the $ADS$ index. The left panel shows that the pricing kernel is a steeper (flatter) U-shaped function of market returns when business conditions are worse (better) than average, as average call returns decrease with the strike and are more (less) negative. As for the right panel, average put returns are more (less) negative when business conditions are better (worse) than average. This resonates the result for volatility, in the sense that the SDF projection is a steeper monotonically decreasing function of negative market returns (i.e., compensation for downside risk is higher) during calm market periods.

Figure 11 contains the results conditioned to option net demand after conditioning on high or low $ADS$. Irrespective of better or worse than average business conditions, the pricing kernel is a steeper U-shaped function of market returns when public investors net sell OTM calls, while the evidence provides more support to a monotonically decreasing SDF when these investors net buy OTM calls. The results for puts during worse than average conditions are similar to those without controlling for $ADS$. In contrast, option net demand is not very informative about the shape of the pricing kernel in the region of negative market returns when business conditions are better than average. Again, this resonates the result for low volatility periods.

In sum, most of the effects of $PNBOC$ and $PNBOP$ are robust to controlling for real business conditions, providing further evidence that option net demand contains unique information about the shape of the pricing kernel that cannot be explained by different macroeconomic conditions.

### 5.3. *Option net demand, controls and the $VRP$*

In this subsection, we use the semi-parametric measure $VRP_t$ for the variance risk premium from Bollerslev, Tauchen and Zhou (2009). An observable time series for the variance risk premium allows us to conduct standard predictive regression analysis, providing additional robustness to our findings. Moreover, we can also include a number of controls simultaneously and test whether the effects of net demand remain the same. More specifically, we first regress $VRP_{t+1}$ on $PNBOC_t$, $PNBOP_t$ and the control variables separately to investigate the univariate predictive relations. Then, we run a multivariate predictive regression including all controls at the same time.[29] In addition to the variables in Section 5.1, we also include as controls the S&P 500 return ($S\&P_t$), which Carr and Wu (2009) show is related to variation in the variance risk premium, and the contemporaneous $VRP_t$.

Table 2 presents the predictability results. Focusing first on the univariate regressions, columns (1) and (2) show that both $PNBOC_t$ and $PNBOP_t$ have strong predictive power for the variance risk premium, with statistically significant negative coefficients. Interestingly, the statistical significance and adjusted $R^2$ are both higher for $PNBOC_t$ than for $PNBOP_t$. This suggests that the net demand for OTM calls may play a more important role in explaining the variance risk premium. Among the remaining predictors, only $LTV_t$ has predictive power, where the higher the volatility from large negative jumps, the higher the $VRP_{t+1}$. This is consistent with Bollerslev, Todorov and Xu (2015), who show that $LTV_t$ captures an important share of the variance risk premium related to compensation for jump tail risk. Contemporaneous $VRP_t$ is also a significant predictor of future $VRP_{t+1}$, indicating some persistence in the variance risk premium.

The last column of Table 2 reports the results for a multivariate regression with all the predictors. As can be seen, the predictive power of $PNBOC_t$ and $PNBOP_t$ remains after including the controls, with statistically significant negative coefficients. The statistical significance of $LTV_t$ and $VRP_t$ also survives the inclusion of the controls. Now two additional variables become significant: $S\&P_t$ and $CJN_t$. The positive relation between market returns and the variance risk premium is consistent with Carr and Wu (2009). The positive coefficient for $CJN_t$ indicates that the $VRP_{t+1}$ tends to be lower following periods with tight intermediary constraints. The remaining variables continue to be insignificant.

[29] The predictive regressions use monthly data from January, 1996 to December, 2019, as the $LTV_t$ is only available until 2019.

The adjusted $R^2$ of the multivariate regression is substantially higher compared to each of the univariate regressions, suggesting that the predictors offer complementary information about the variance risk premium.

The significant negative relation between public net demand of OTM options and the variance risk premium is consistent with our theory and the interpretation that the premium for variance risk reflects the compensation for the risk exposures of public investors. Intuitively, when public investors net sell OTM calls and OTM puts, they get exposed to market variance risk and require more compensation to bear it, increasing the variance risk premium. Conversely, when public investors net buy OTM calls and OTM puts, their exposure to variance risk decreases, as does the premium to bear it.

As a final test, we also analyze how the average premium for variance risk changes depending on the net demand in the option market. This analysis is similar to that in Section 4.2, but using directly the $VRP$ measure. More specifically, Figure 12 reports the average of the $VRP_{t+1}$ calculated in the months $t+1$ following the months $t$ where net demand is below or above several thresholds. Over the whole sample, the average variance risk premium is 1.73%. As the left panel of the figure shows, the average $VRP_{t+1}$ is higher when public investors are exposed to market upside risk by net selling OTM calls, and considerably lower when they are net buying these options. In fact, the average premium can be even negative (although not statistically different from zero) when $PNBOC$ is very large. The middle panel displays a similar pattern when conditioning on $PNBOP$: the average $VRP_{t+1}$ is higher when public investors are more exposed to market downside risk by net selling OTM puts, and much lower when they are net buying these options.

Our theory predicts that exposure to market variance risk, and compensation for that risk, should be even higher (lower) when public net demand is negative (positive) for both OTM calls and OTM puts at the same time. The right panel of Figure 12 reports the average variance risk premium when both $PNBOC$ and $PNBOP$ are below or above several fractions of their respective standard deviations. As expected, the differences in the average $VRP_{t+1}$ are even more extreme in this case (albeit with less statistical precision due to smaller samples). In particular, when public investors are net selling large quantities of OTM options, the average premium is almost twice as high as for the whole sample. In contrast, the average $VRP_{t+1}$ is highly negative when public investors are net buying large quantities of both OTM calls and OTM puts.

The relation between demand in the option market and the $VRP$ has been previously investigated in the literature in different ways. Garleanu, Pedersen and Poteshman (2009) find that public net demand for all options is positively related to the variance risk premium. Their interpretation is that demand pressure increases the relative expensiveness of the options, as measured by the $VRP$. With a longer sample period, Chen, Joslin and Ni (2019) document a negative relation between different measures of net demand and $VRP$, arguing that supply shocks to market makers, which mostly occur when they are financially constrained, have a predominant effect on option expensiveness over demand shocks. Moreover, Fournier and Jacobs (2020) show that the variance risk premium is higher when market makers' inventory risk is high, which happens when the accumulation of past net-buying volumes of public investors for all options is positive and large.

Our findings are novel as we distinguish between net demand for OTM calls and OTM puts and show that both negatively predict the $VRP$. This negative relation cannot be explained by demand pressure (as in this case the sign should be positive) and is robust to measures of intermediary and funding constraints. Therefore, the empirical evidence supports that the negative relation results from the compensation required by public investors for variance risk arising from the exposure of new OTM options positions. In a broad sense, this is consistent with Chen, Joslin and Ni (2019) and Fournier and Jacobs (2020). When market makers face tight financial constraints and high inventory risk, they need to buy OTM options to decrease their inventories' variance risk. Public investors are only willing to sell these options at relatively high prices to compensate for the exposure to variance risk of these new positions, thus increasing the $VRP$.

# 6. Conclusion

The pricing kernel recovered from option prices is usually locally increasing in market returns, which contradicts risk averse behavior in a representative investor framework. Call and put index option returns are also too low to be consistent with standard asset pricing theory. Relatedly, the observed premium for bearing variance risk is considerably high. In this paper, we connect these puzzles to net demand in the option market. If options are nonredundant, trading quantities in the option market should be informative about the risk exposures, and, consequently, compensation for risk of the marginal investor in this market.

Empirically, we document strikingly different patterns in average option returns and, consequently, in the shape of the pricing kernel, depending on the option quantities that public investors net buy or net sell in the S&P 500 option market. We also show that there is a strong negative relation between option net demand for OTM options and the variance risk premium. These novel empirical patterns are consistent with our theoretical framework and the idea that public investors can be seen as the marginal investor in the option market. More generally, our empirical evidence supports heterogeneous-agent equilibrium models in which options are nonredundant.

# A. Proofs

Since the option returns are functions $f(R, K_j)$ of the market returns $R$, we can write $\mathbf{R}^e = \mathbf{g}(R)$, where $\mathbf{g}(R) = [f(R, K_1), f(R, K_2), ..., f(R, K_{N-1}), R]^{\mathrm{T}} - R_f \mathbf{1}_N$, such that $\mathbf{g}'(R) = [f'(R, K_1), f'(R, K_2), ..., f'(R, K_{N-1}), 1]^{\mathrm{T}}$. Therefore, let $m^*(R) = U'(W_0 R_f + \lambda^* \mathbf{g}(R))$ be the option-implied pricing kernel. Let also: $\lambda_N$ denote the portfolio weight for the market; $\lambda_j$, $j = 1, ..., N_c$, denote the weights for the OTM calls, where $K_j/S_t > 1$; and $\lambda_j$, $j = N_c + 1, ..., N - 1$, denote the weights for the OTM puts, where $K_j/S_t < 1$.

## *A.1. Proposition 1*

We consider that the investor associated with the pricing kernel is long in the market index, that is, $\lambda_N^* > 0$. We now assume, by contradiction, that the pricing kernel is a U-shaped function of $R$ and that the investor is long in all OTM calls. More precisely, we assume that $m^{*\prime}(R_u) > 0$ at an upper market return $R_u > 1$ and that $\lambda_j^* > 0$ for $j = 1, ..., N_c$. Taking the derivative of $m^*(R_u)$ with respect to $R_u$, we obtain:

$$m^{*\prime}(R_u) = U''(W_0 R_f + \lambda^* \mathbf{g}(R_u))\, \lambda^* \mathbf{g}'(R_u). \tag{A.1}$$

First, note that since $U(.)$ is a concave function, $U''(W_0 R_f + \lambda^* \mathbf{g}(R_u)) < 0$. Second, for $R_u > 1$, we have for all OTM put options that $f(R_u, K_j) = 0$, and, hence, that $f'(R_u, K_j) = 0$. For OTM call options, we have $f'(R_u, K_j) > 0$ for $K_j/S_t < R_u$ and $f'(R_u, K_j) = 0$ for $K_j/S_t > R_u$. Therefore, $\mathbf{g}'(R_u) \geq 0$. Finally, if $\lambda_j^* > 0$ for all $j = 1, ..., N_c$ OTM calls, we have $\lambda^* \mathbf{g}'(R_u) > 0$ and, hence, $m^{*\prime}(R_u) < 0$, which contradicts the initial assumption. To have $m^{*\prime}(R_u) > 0$, we need at least one OTM call with $\lambda_j^* < 0$.

## *A.2. Proposition 2*

We consider that the investor associated with the pricing kernel is long in the market index, that is, $\lambda_N^* > 0$. We now assume, by contradiction, that the pricing kernel is an increasing function of $R$ at a low market return and that the investor is short in all OTM puts. More precisely, we assume that $m^{*\prime}(R_d) > 0$ at a low market return $R_d < 1$ and that $\lambda_j^* < 0$ for $j = N_c + 1, ..., N - 1$. Taking the derivative of $m^*(R_d)$ with respect to $R_d$, we obtain:

$$m^{*\prime}(R_d) = U''(W_0 R_f + \lambda^* \mathbf{g}(R_d))\, \lambda^* \mathbf{g}'(R_d). \tag{A.2}$$

First, note that since $U(.)$ is a concave function, $U''(W_0 R_f + \lambda^* \mathbf{g}(R_d)) < 0$. Second, for $R_d < 1$, we have for all OTM call options that $f(R_d, K_j) = 0$, and, hence, that $f'(R_d, K_j) = 0$. For OTM put options, we have $f'(R_d, K_j) < 0$ for $K_j/S_t > R_d$ and $f'(R_d, K_j) = 0$ for $K_j/S_t < R_d$. Therefore, $\mathbf{g}'(R_d) \leq 0$. Finally, if $\lambda_j^* < 0$ for all $j = N_c + 1, ..., N - 1$ OTM puts, we have $\lambda^* \mathbf{g}'(R_d) > 0$ and, hence, $m^{*\prime}(R_d) < 0$, which contradicts the initial assumption. To have $m^{*\prime}(R_d) > 0$, we need at least one OTM put with $\lambda_j^* > 0$.

# B. Pricing kernel estimation and marginal investor interpretation

Consider the excess returns $\mathbf{R}^e$ of the market index and the index options defined in Section 2.2. In this setting, an admissible SDF is a random variable $m$ such that:

$$\mathbb{E}[m\mathbf{R}^e] \equiv \int m\mathbf{R}^e \,\mathrm{d}\mathbb{P} = \mathbf{0}_N, \tag{B.1}$$

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

$$\int \mathbf{R}^e \frac{m}{\mathbb{E}[m]}\,\mathrm{d}\mathbb{P} = \int \mathbf{R}^e \,\mathrm{d}\mathbb{Q} \equiv \mathbb{E}^{\mathbb{Q}}[\mathbf{R}^e] = \mathbf{0}_N. \tag{B.2}$$

Under the restriction that $\mathbb{Q}$ is absolutely continuous with respect to $\mathbb{P}$, $\mathrm{d}\mathbb{Q}/\mathrm{d}\mathbb{P}$ is a Radon-Nikodym derivative.[30]

Consider now the problem of estimating the risk-neutral measure that minimizes a general convex loss function $\phi$ with respect to the physical measure, while correctly pricing the market index and options:

$$\mathbb{Q}^* = \arg\min_{Q} \mathbb{E}\left[\phi\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)\right] \equiv \int \phi\left(\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right)\mathrm{d}\mathbb{P}, \text{ s.t. } \mathbb{E}^{\mathbb{Q}}\left[\mathbf{R}^e\right] = \mathbf{0}_N, \tag{B.3}$$

where the risk-neutral measure must also be nonnegative and integrate to one. Rubinstein

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

(1994), Jackwerth and Rubinstein (1996) and Stutzer (1996) consider similar problems for specific loss functions. In the absence of arbitrage, problem (B.3) is equivalent to the simpler dual problem below (Almeida and Garcia, 2017; Almeida and Freire, 2022):

$$\underset{\lambda \in \Lambda}{max}\ \mathbb{E}\left[-\phi_{+}^{*}(1-\lambda \mathbf{R}^{e})\right], \tag{B.4}$$

where $\phi_{+}^{*}$ denotes the convex conjugate of $\phi$ and $\Lambda = \{\lambda \in \mathbb{R}^{N} : (1-\lambda \mathbf{R}^{e}) \in \text{dom } \phi_{+}^{*}\}$.[31] More precisely, the solution to (B.3) is obtained from the first order condition of (B.4):

$$\frac{\mathrm{d}\mathbb{Q}^{*}(\mathbf{R})}{\mathrm{d}\mathbb{P}} = (\phi_{+}^{*})'(1-\lambda^{*}\mathbf{R}^{e}), \tag{B.5}$$

where the row-vector $\lambda$ is composed of the Lagrange multipliers coming from the Euler equations in (B.3) for the basis assets, i.e., the market index and the index options. That is, the assets that enter the pricing kernel are those that the SDF is required to price. The pricing kernel is recovered as $m^{*}(\mathbf{R}) = \frac{1}{R_f}\,(\phi_{+}^{*})'(1-\lambda^{*}\mathbf{R}^{e})$.

The dual problem (B.4) can be economically interpreted as a standard optimal portfolio problem for an investor maximizing a general concave utility function ($u = -\phi_{+}^{*}$) over future wealth, where $\lambda$ are the portfolio weights for the market index and the options.[32] In this sense, $m^{*}(\mathbf{R})$ is monotonically decreasing in the optimal portfolio returns $\lambda^{*}\mathbf{R}^{e}$ and proportional to the marginal utility of the risk averse investor. From the first order condition of the dual problem, the Euler equations for the market index and the index options are satisfied, such that the investor is marginal in the index and option markets.

A distinctive feature of the dual problem solution is that it identifies the pricing kernel as a function of the basis assets returns. As an implication, one can project $m^{*}(\mathbf{R})$ onto the returns of any of the basis assets, and, in particular, of the market index. This provides an interpretation for the standard approach in (3): $m(R)$ is the projection of $m^{*}(\mathbf{R})$ onto market return states. Importantly, given a physical distribution, this interpretation is valid in population regardless of the approach used to estimate $m(R)$ due to the uniqueness of the risk-neutral distribution and, hence, of the pricing kernel projection, for a sufficiently large cross-section of options.[33]

[31] The convex conjugate is given by $\phi_{+}^{*}(z) = \underset{w \in [0,\infty) \cap \text{dom } \phi}{sup} zw - \phi(w)$, while the domain of $\phi_{+}^{*}(z)$ is defined as the values of $z$ for which the function is finite ($\phi_{+}^{*}(z) < \infty$).

[32] In particular, if the loss function $\phi$ belongs to the Cressie and Read (1984) family, Almeida and Freire (2022) show that the dual problem is equivalent to the maximization of a hyperbolic absolute risk aversion (HARA) utility function.

[33] The spanning properties of option prices also imply that the specific convex loss function $\phi$ in (B.3)

# C. Shape of the pricing kernel and expected option returns

The shape of the pricing kernel is directly related to the expected returns of call and put options, $\mu_t^c$ and $\mu_t^p$, respectively. Expected option returns can be defined as below:[34]

$$\mu_t^c(S_T, K) = \frac{\mathbb{E}_t[(S_T - K)^+]}{\mathbb{E}_t[(S_T - K)^+ \, m_{t,T}(S_T)]} - 1, \tag{C.1}$$

$$\mu_t^p(S_T, K) = \frac{\mathbb{E}_t[(K - S_T)^+]}{\mathbb{E}_t[(K - S_T)^+ \, m_{t,T}(S_T)]} - 1, \tag{C.2}$$

where $x^+ = \max(x, 0)$, $K$ is the strike price and $S_T = S_t R_{t,T}$ is the value for the market index at time $T$.[35] The numerator in (C.1) and (C.2) is the expected payoff of the option under the physical measure, while the denominator is the expected payoff under the risk-neutral measure, i.e., the option price.

Coval and Shumway (2001) show that if $m_{t,T}(S_T)$ is monotonically decreasing, both calls and puts have expected returns that increase with the strike price.[36] Intuitively, a monotonically decreasing $m_{t,T}(S_T)$ shifts probability mass towards states where the call (put) option is less (more) valuable. Therefore, as the strike increases, the call price decreases by more than the expected payoff, increasing the expected return. Conversely, as the strike decreases, the put price decreases by less than the expected payoff, decreasing the expected return. They also demonstrate that expected call (put) returns are positive (negative) under a monotonically decreasing pricing kernel.[37]

Bakshi, Madan and Panayotov (2010) consider the alternative scenario where the pricing kernel is a U-shaped function of market returns.[38] In this case, they show that expected call returns are decreasing in the strike and negative beyond a strike threshold, where the steeper the slope of the SDF in its increasing region, the more negative are the expected returns. The intuition is that $m_{t,T}(S_T)$ shifts probability mass towards states where the call option is more valuable, such that the call price decreases by less than

is unimportant, as the pricing restrictions for the options completely determine the solution.

[34] To see that, note that $\mathbb{E}_t\left[\frac{(S_T-K)^+}{\mathbb{E}_t[(S_T-K)^+ \, m_{t,T}(S_T)]}\right] = \frac{\mathbb{E}_t[(S_T-K)^+]}{\mathbb{E}_t[(S_T-K)^+ \, m_{t,T}(S_T)]}$.

[35] Note that it is equivalent to denote the projected pricing kernel as $m_{t,T}(S_T)$ or $m_{t,T}(R_{t,T})$.

[36] Chaudhuri and Schroder (2015) generalize this result by showing that the expected returns of payoffs in the log-concave class are increasing in the strike under a monotonically decreasing pricing kernel.

[37] More precisely, they show that expected put returns should be below the risk-free rate, but in practice expected put returns are virtually always negative if the SDF is monotonically decreasing.

[38] That is, there exists an upper market return $R_u$ such that $m'_{t,T}(R_{t,T}) > 0$ and $m''_{t,T}(R_{t,T}) > 0$ for $R_{t,T} > R_u$, where $R_u \geq R_d$ and $m_{t,T}$ is monotonically decreasing for $R_{t,T} < R_d$.

the expected payoff as the strike increases. This decreases the expected returns, which eventually reach negative values. Since a U-shaped SDF is still declining in the region of negative market returns that is relevant for put options, the implications for expected put returns are the same as under a monotonically decreasing SDF.

To illustrate the relations above, we construct hypothetical pricing kernels with different shapes and show their implications for expected option returns. More specifically, we take as the hypothetical physical measure the unconditional physical distribution of 30-day market returns over our sample period, and specify pricing kernels as distinct functions of those market returns. Figure C.1 reports the results. In the upper panels, we have the implications of monotonically decreasing pricing kernels with different slopes. Consistent with Coval and Shumway (2001), expected call (put) returns are positive (negative) and increase with the strike. In particular, the steeper the pricing kernel, the higher (more negative) are the expected call (put) returns.

The middle panels of Figure C.1 illustrate the patterns for monotonically increasing pricing kernels with different slopes. In this case, implications are the opposite of those of a monotonically decreasing SDF. Expected call (put) returns are negative (positive) and decrease with the strike. This is because the pricing kernel shifts probability mass towards states where the call (put) option is more (less) valuable. The steeper the pricing kernel, the more negative (higher) are the expected call (put) returns. The lower panels depict the implications of a U-shaped SDF. Aligned with Bakshi, Madan and Panayotov (2010), expected call returns are negative and decrease with the strike. On the other hand, since the SDF is still monotonically decreasing for negative (net) market returns, expected put returns are negative and increase with the strike.

The relations described above hold for expected option returns and the pricing kernel conditional on time $t$. Working with these objects empirically is challenging as the conditional physical distribution $d\mathbb{P}_t$ is required to compute $\mathbb{E}_t$. There is no consensus on how to estimate $d\mathbb{P}_t$ such that it reflects the same conditioning information set that investors have available when setting option prices (see, e.g., the discussion in Linn, Shive and Shumway, 2017). To circumvent this issue, Coval and Shumway (2001), Bakshi, Madan and Panayotov (2010) and Chaudhuri and Schroder (2015) condition down by taking unconditional expectations of option returns, which can be directly obtained by computing averages of realized option returns.

As discussed by Chaudhuri and Schroder (2015), a rejection of increasing monotonicity of expected returns across strikes based on unconditional returns implies rejection with conditional returns (if the unconditional expectation of return differences across strikes is negative, the conditional expectation cannot be a.s. positive). In other words, a rejection based on unconditional expected returns suffices to identify that there is nonmonotonicity of the pricing kernel in at least one date (and likely on most dates) of the sample. On the other hand, failure to reject increasing monotonicity of expected returns across strikes based on unconditional returns does not imply that for every $t$ the conditional pricing kernel is monotonically decreasing. Rather, it indicates that at least for most dates the pricing kernel is monotonically decreasing, i.e., that this is the overall shape in the sample.

# D. Option net demand and expected option returns

Table D.1 reports the average call option returns calculated over the whole sample and conditioned to different months based on $PNBOC$. The observed values are the same as those depicted in the left panel of Figure 2, with the difference that in the figure the deltas are replaced by the average moneyness of the options for a given delta. The table also presents the average difference between the returns of the option in the current column and the previous one. From left to right, we observe average returns of calls with increasing strike, going from ITM to OTM. We report whether values are statistically different from zero based on bootstrap draws from the respective sample of option returns.

The results for the whole sample are in line with the previous literature (Bakshi, Madan and Panayotov, 2010; Chaudhuri and Schroder, 2015; Baele et al., 2019). Average call option returns are positive for ITM calls, but decrease as the strike increases until they become negative for OTM calls. The decrease of the average return with respect to the strike is statistically significant from delta equal to 0.50 to lower deltas. This evidence rejects the monotonicity of the pricing kernel for the whole sample, supporting an overall U-shaped SDF projection.

However, the patterns in average call returns are remarkably different once we condition to months where public investors are net buyers or net sellers of OTM calls. When $PNBOC > 0$, average returns are always positive and increase with the strike. While there is a small decrease in the average return between deltas 0.20 and 0.25, the difference is not statistically different from zero. In contrast, when $PNBOC < 0$, average returns are substantially smaller than for the whole sample and decrease with the strike at a faster rate, becoming significantly negative from delta 0.45 onwards. The average return difference is always negative and significant, i.e., average returns decrease with the strike.

The findings above suggest that the evidence from unconditional expected call returns calculated in the whole sample masks the different patterns in call returns when public investors are net buying or net selling OTM calls. When $PNBOC > 0$, average call returns are consistent with a monotonically decreasing projection of the pricing kernel in the region of positive net market returns. On the other hand, when public investors are net selling OTM calls, the average returns of call options provide much stronger evidence of a U-shaped pricing kernel than that obtained from unconditional average returns.

The different patterns in average call returns are even more striking once we consider months where $PNBOC >$ 1std and $PNBOC < -$1std. When public investors are net buying large amounts of OTM calls, average returns are highly positive and increase with the strike at a fast rate, where all average return differences are positive and statistically significant. This indicates a steep monotonically decreasing projection of the SDF in the market upside. On the other hand, in months where public investors net sell large amounts of OTM calls, the average call returns are highly negative and decrease significantly with the strike, supporting a steeper U-shaped SDF projection.

In Table D.2, we report the average put option returns calculated over the whole sample and conditioned to different months based on $PNBOP$. Observed values are the same as those depicted in the right panel of Figure 2. From left to right, we observe average returns of puts with increasing strike, going from OTM to ITM. Again, the results for the whole sample are in line with the previous literature. Average put option returns are highly negative and increase with the strike. This indicates that the SDF is monotonically decreasing in the region of negative net market returns, with a steep slope.

Once we condition to months where $PNBOP > 0$ ($PNBOP < 0$), average put returns are less (more) negative and increase with the strike at a slower (faster) rate, consistent with a flatter (steeper) decreasing SDF projection. The SDF projection gets even steeper when $PNBOP < -$1std, with extremely negative average put returns. However, a strikingly different pattern arises when $PNBOP >$ 1std. In this case, average put returns are highly positive and decreasing in the strike. With the exception of the options with deltas from $-0.70$ to $-0.80$, all average returns are statistically different from zero. The decrease of the average return with respect to the strike is statistically significant for all cases. This rejects the decreasing monotonicity of the pricing kernel in the region of negative net market returns, supporting an increasing SDF projection.

# References


Adrian, T., Etula, E., Muir, T., 2014. Financial intermediaries and the cross-section of asset returns. The Journal of Finance 69, 2557-2596.

Ait-Sahalia, Y., Lo, A. W., 2000. Nonparametric risk management and implied risk aversion. Journal of Econometrics 94, 9-51.

Ait-Sahalia, Y., Wang, Y., Yared, F., 2001. Do option markets correctly price the probabilities of movement of the underlying asset? Journal of Econometrics 102, 67-110.

Almeida, C., Garcia, R., 2017. Economic implications of nonlinear pricing kernels. Management Science 63, 3147-3529.

Almeida, C., Freire, G., 2022. Pricing of index options in incomplete markets. Journal of Financial Economics 144, 174-205.

Andersen, T. G., Fusari, N., Todorov, V., 2015. The risk premia embedded in index options. Journal of Financial Economics 117, 558-584.

Aruoba, S. B., Diebold, F. X., Scotti, C., 2009. Real-time measurement of business conditions. Journal of Business and Economic Statistics 27, 417-427.

Baele, L., Driessen, J., Ebert, S., Londono, J. M., Spalt, O. G., 2019. Cumulative prospect theory, option returns, and the variance risk premium. The Review of Financial Studies 32, 3667-3723.

Bakshi, G., Cao, C., Chen, Z., 2000. Do call prices and the underlying stock always move in the same direction? The Review of Financial Studies 13, 549–584.

Bakshi, G., Kapadia, N., 2003. Delta-hedged gains and the negative market volatility risk premium. The Review of Financial Studies 16, 527–566.

Bakshi, G., Madan, D., Panayotov, G., 2010. Returns of claims on the upside and the viability of u-shaped pricing kernels. Journal of Financial Economics 97, 130–154.

Barone-Adesi, G., Engle, R., Mancini, L., 2008. A GARCH option pricing model with filtered historical simulation. The Review of Financial Studies 21, 1223–1258.

Barone-Adesi, G., Fusari, N., Mira, A., Sala, C., 2020. Option market trading activity and the estimation of the pricing kernel: a Bayesian approach. Journal of Econometrics 216, 430-449.

Bates, D., 2000. Post-'87 crash fears in the S&P 500 futures option market. Journal of Econometrics 94, 181–238.

Bates, D., 2008. The market for crash risk. Journal of Economic Dynamics and Control 32, 2291–2321.

Beare, B. K., Schmidt, L., 2016. An empirical test of pricing kernel monotonicity. Journal of Applied Econometrics 31, 338–356.

Beason, T., Schreindorfer, D., 2022. Dissecting the Equity Premium. Journal of Political Economy 130, 1993-2225.

Bhamra, H. S., Uppal, R., 2009. The effect of introducing a non-redundant derivative on the volatility of stock-market returns when agents differ in risk aversion. The Review of Financial Studies 22, 2303–2330.

Bollen, N. P. B., Whaley, R., 2004. Does net buying pressure affect the shape of implied volatility functions? The Journal of Finance 59, 711-753.

Bollerslev, T., Tauchen, G., Zhou, H., 2009. Expected stock returns and variance risk premia. The Review of Financial Studies 22, 4463–4492.

Bollerslev, T., Todorov, V., Xu, L., 2015. Tail risk premia and return predictability. Journal of Financial Economics 118, 113-134.

Bondarenko, O., 2014. Why are put options so expensive? The Quarterly Journal of Finance 4, 145-195.

Breeden, D., Litzenberger, R. H., 1978. Prices of state-contingent claims implicit in option prices. Journal of Business 51, 621-651.

Brunnermeier, M. K., Sannikov, Y., 2014. A macroeconomic model with a financial sector. American Economic Review 104, 379-421.

Brown, D., Jackwerth, J. C., 2012. The pricing kernel puzzle: reconciling index option data and economic theory. In: Contemporary Studies in Economics and Financial Analysis: Derivative Securities Pricing and Modelling 94, 155–183. Bingley: Emerald Group.

Buraschi, A., Jackwerth, J. C., 2001. The price of a smile: Hedging and spanning in option markets. The Review of Financial Studies 14, 495–527.

Buraschi, A., Jiltsov, A., 2006. Model uncertainty and option markets with heterogeneous beliefs. The Journal of Finance 61, 2841-2897.

Carr, P., Wu, L., 2009. Variance risk premiums. The Review of Financial Studies 22, 1311–1341.

Chabi-Yo, F., Garcia, R., Renault, E., 2008. State dependence can explain the risk aversion puzzle. The Review of Financial Studies 21, 973–1011.

Chaudhuri, R., Schroder, M., 2015. Monotonicity of the stochastic discount factor and expected option returns. The Review of Financial Studies 28, 1462-1505.

Chen, H., Joslin, S., Ni, S., 2019. Demand for crash insurance, intermediary constraints and risk premia in financial markets. The Review of Financial Studies 32, 228-265.

Christoffersen, P., Heston, S., Jacobs, K., 2013. Capturing option anomalies with a variance-dependent pricing kernel. The Review of Financial Studies 26, 1962–2006.

Cochrane, J., 2001. Asset pricing. Princeton University Press, Princeton.

Coval, J. D., Shumway, T., 2001. Expected option returns. The Journal of Finance 56, 983-1009.

Cressie, N., Read, T., 1984. Multinomial goodness-of-fit tests. Journal of the Royal Statistical Society Series B (Methodological) 46, 440-464.

Cuesdeanu, H., Jackwerth, J. C., 2018. The pricing kernel puzzle: survey and outlook. Annals of Finance 14, 289-329.

Driessen, J., Maenhout, P., 2007. An empirical perspective on option pricing anomalies. Review of Finance 11, 561-603.

Farago, A., Khapko, M., Ornthanalai, C., 2021. Asymmetries and the market for put options. Working paper, University of Toronto.

Fournier, M., Jacobs, K., 2020. A tractable framework for option pricing with dynamic market maker inventory and wealth. Journal of Financial and Quantitative Analysis 55, 1117-1152.

Franke, G., Stapleton, R. C., Subrahmanyam, M. G., 1998. Who buys and who sells options: The role of options in an economy with background risk. Journal of Economic Theory 82, 89-109.

Gabaix, X., Krishnamurthy, A., Vigneron, O., 2007. Limits of arbitrage: Theory and evidence from the mortgage-backed securities market. The Journal of Finance 62, 557-595.

Gârleanu, N., Pedersen, L. H., Poteshman, A. M., 2009. Demand-based option pricing. The Review of Financial Studies 22, 4259–4299.

Grith, M., Hardle, W. K., Kratschmer, V., 2017. Reference-dependent preferences and the empirical pricing kernel puzzle. Review of Finance 21, 269-298.

Gruenthaler, T., 2022. Risk premia and option intermediation. Working paper.

He, Z., Kelly, B., Manela, A., 2017. Intermediary asset pricing: New evidence from many asset classes. Journal of Financial Economics 126, 1-35.

He, Z., Krishnamurthy, A., 2013. Intermediary asset pricing. American Economic Review 103, 732-770.

Hens, T., Reichlin, C., 2013. Three solutions to the pricing kernel puzzle. Review of Finance 17, 1065–1098.

Heston, S. L., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. The Review of Financial Studies 6, 327–343.

Heston, S. L., Nandi, S., 2000. A closed-form GARCH option valuation model. The Review of Financial Studies 13, 585–625.

Jacobs, K., Mai, A. T., Pederzoli, P., 2022. Identifying demand and supply in index option markets. Working paper.

Jackwerth, J. C., Rubinstein, M., 1996. Recovering probability distributions from option prices. The Journal of Finance 51, 1611-1631.

Jackwerth, J. C., 2000. Recovering risk aversion from option prices and realized returns. The Review of Financial Studies 13, 433–451.

Johnson, T., Liang, M., Liu, Y., 2018. What drives index options exposures? Review of Finance 22, 561–593.

Jones, C., 2006. A nonlinear factor analysis of S&P 500 index option returns. The Journal of Finance 61, 2325–2363.

Lemmon, M., Ni, S., 2014. Differences in trading and pricing between stock and index options. Management Science 60, 1985-2001.

Linn, M., Shive, S., Shumway, T., 2018. Pricing kernel monotonicity and conditional information. The Review of Financial Studies 31, 493-531.

Lucas, R., 1978. Asset prices in an exchange economy. Econometrica 46, 1429–1445.

Muravyev, D., 2016. Order flow and expected option returns. The Journal of Finance 71, 673-707.

Polkovnichenko, V., Zhao, F., 2013. Probability weighting functions implied in option prices. Journal of Financial Economics 107, 580-609.

Rosenberg, J., Engle, R., 2002. Empirical pricing kernels. Journal of Financial Economics 64, 341-372.

Rubinstein, M., 1994. Implied binomial trees. The Journal of Finance 49, 771–818.

Santa-Clara, P., Saretto, A., 2009. Option strategies: Good deals and margin calls. Journal of Financial Markets 12, 391-417.

Schreindorfer, D., Sichert, T., 2022. Volatility and the pricing kernel. Working paper.

Sichert, T., 2022. The pricing kernel is U-shaped. Working paper.

Song, Z., Xiu, D., 2016. A tale of two option markets: state-price densities and volatility risk. Journal of Econometrics 190, 176–196.

Stutzer, M., 1996. A simple nonparametric approach to derivative security valuation. The Journal of Finance 51, 1633–1652.

Vanden, J. M., 2004. Options trading and the CAPM. The Review of Financial Studies 17, 207–238.

Vanden, J. M., 2006. Option coskewness and capital asset pricing. The Review of Financial Studies 19, 1279–1320.

Welch, I., Goyal, A., 2008. A comprehensive look at the empirical performance of equity premium prediction. The Review of Financial Studies 21, 1455–1508.

Ziegler, A., 2007. Why does implied risk aversion smile? The Review of Financial Studies 20, 859–904.

Table 1: **Option net demand and alternative determinants**

| | Dependent Variable: $PNBOC_t$ | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | $D_t^{NBER}$ | $RV_t$ | $LTV_t$ | $TMS_t$ | $CRS_t$ | $DFS_t$ | $ADS_t$ | $CJN_t$ | $TED_t$ |
| Slope | −0.18*** | −0.28*** | −0.13** | −0.28*** | −0.13** | −0.22*** | 0.02 | 0.20*** | −0.18*** |
| | (−3.25) | (−5.04) | (−2.24) | (−5.07) | (−2.21) | (−3.93) | (0.36) | (3.52) | (−3.26) |
| Adj. $R^2$ (%) | 3.10 | 7.56 | 1.38 | 7.64 | 1.29 | 4.63 | −0.29 | 3.68 | 3.13 |
| | Dependent Variable: $PNBOP_t$ | | | | | | | | |
| | $D_t^{NBER}$ | $RV_t$ | $LTV_t$ | $TMS_t$ | $CRS_t$ | $DFS_t$ | $ADS_t$ | $CJN_t$ | $TED_t$ |
| Slope | 0.11* | 0.16*** | 0.01 | −0.36*** | 0.01 | −0.01 | −0.18*** | 0.44*** | 0.17*** |
| | (1.94) | (2.82) | (0.25) | (−6.62) | (0.02) | (−0.06) | (−3.24) | (8.65) | (3.13) |
| Adj. $R^2$ (%) | 0.92 | 2.28 | −0.33 | 12.56 | −0.34 | −0.33 | 3.09 | 19.83 | 2.87 |

This table presents the results for univariate monthly regressions of $PNBOC_t$ ($PNBOP_t$) on a set of regressors indicated by the columns. Variables are standardized to have unit variance to facilitate interpretation. Each row reports the estimated slope coefficient and, in parenthesis, $t$-statistics, with the exception of the last row which depicts the adjusted $R^2$. The intercept is included in the regressions but is not reported in the table for brevity. *, ** and *** represent statistical significance at the 10%, 5% and 1% levels, respectively. The sample ranges from January, 1996 to December, 2020, with the exception of the regression on $LTV_t$ which ends at December, 2019.

Table 2: **Variance risk premium predictability**

| | Dependent Variable: $VRP_{t+1}$ | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | (11) | (12) | (13) |
| Intercept | 0.72*** | 0.71*** | 0.70*** | 0.54** | 0.12 | 0.61*** | 0.65** | 0.62** | 0.74*** | 0.71*** | 0.79*** | 0.53*** | 0.08 |
| | (10.91) | (11.18) | (9.16) | (2.22) | (0.60) | (6.12) | (2.32) | (2.46) | (13.95) | (10.80) | (4.69) | (6.11) | (0.30) |
| $PNBOC_t$ | −0.19*** | | | | | | | | | | | | −0.15** |
| | (−4.81) | | | | | | | | | | | | (−2.14) |
| $PNBOP_t$ | | −0.14*** | | | | | | | | | | | −0.08** |
| | | (−2.66) | | | | | | | | | | | (−2.00) |
| $S\&P_t$ | | | 0.09 | | | | | | | | | | 0.17*** |
| | | | (0.78) | | | | | | | | | | (2.72) |
| $RV_t$ | | | | 0.09 | | | | | | | | | 0.18 |
| | | | | (0.56) | | | | | | | | | (1.03) |
| $LTV_t$ | | | | | 0.23*** | | | | | | | | 0.29** |
| | | | | | (2.80) | | | | | | | | (1.98) |
| $TMS_t$ | | | | | | 0.06 | | | | | | | −0.12 |
| | | | | | | (0.92) | | | | | | | (−1.19) |
| $CRS_t$ | | | | | | | 0.02 | | | | | | −0.11 |
| | | | | | | | (0.19) | | | | | | (−0.88) |
| $DFS_t$ | | | | | | | | 0.04 | | | | | 0.05 |
| | | | | | | | | (0.33) | | | | | (0.37) |
| $ADS_t$ | | | | | | | | | 0.12 | | | | 0.14 |
| | | | | | | | | | (0.64) | | | | (0.77) |
| $CJN_t$ | | | | | | | | | | 0.01 | | | 0.07** |
| | | | | | | | | | | (0.28) | | | (2.13) |
| $TED_t$ | | | | | | | | | | | −0.06 | | −0.14 |
| | | | | | | | | | | | (−0.38) | | (−1.17) |
| $VRP_t$ | | | | | | | | | | | | 0.25** | 0.21** |
| | | | | | | | | | | | | (2.55) | (2.38) |
| Adj. $R^2$ (%) | 3.13 | 1.72 | 0.38 | 0.46 | 5.04 | −0.02 | −0.31 | −0.20 | 1.20 | −0.34 | 0.04 | 6.18 | 17.98 |

This table presents the results for monthly predictive regressions of the variance risk premium on a set of predictors indicated by the rows. Columns (1) to (12) contain the results for univariate regressions, while column (13) contains the results for a multivariate regression including all the predictors. Variables are standardized to have unit variance to facilitate interpretation. Each row reports the estimated coefficient and, in parenthesis, Newey-West $t$-statistics with a lag length equal to one, with the exception of the last row which depicts the adjusted $R^2$. *, ** and *** represent statistical significance at the 10%, 5% and 1% levels, respectively. The sample ranges from January, 1996 to December, 2019.

Table D.1: **Option net demand and average call option returns**

| Delta | 0.80 | 0.75 | 0.70 | 0.65 | 0.60 | 0.55 | 0.50 | 0.45 | 0.40 | 0.35 | 0.30 | 0.25 | 0.20 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | Whole Sample | | | | | | |
| Avg. Ret. | 4.46*** | 4.59*** | 4.75*** | 4.72*** | 4.65*** | 4.41*** | 3.80** | 2.94* | 1.80 | 0.05 | −2.32 | −5.65** | −10.70*** |
| Avg. Diff. | | 0.13 | 0.15 | −0.03 | −0.06 | −0.25 | −0.61*** | −0.86*** | −1.15*** | −1.75*** | −2.37*** | −3.33*** | −5.06*** |
| | | | | | | | $PNBOC > 0$ | | | | | | |
| Avg. Ret. | 5.99*** | 6.90*** | 7.93*** | 8.78*** | 9.62*** | 10.48*** | 11.23*** | 12.06*** | 12.88*** | 13.70*** | 14.63*** | 15.29*** | 14.68*** |
| Avg. Diff. | | 0.92*** | 1.03*** | 0.85*** | 0.84*** | 0.86*** | 0.76*** | 0.82** | 0.82** | 0.82* | 0.93* | 0.66 | −0.61 |
| | | | | | | | $PNBOC < 0$ | | | | | | |
| Avg. Ret. | 3.07*** | 2.49** | 1.84 | 1.00 | 0.12 | −1.13 | −2.98 | −5.37*** | −8.31*** | −12.40*** | −17.79*** | −24.75*** | −33.87*** |
| Avg. Diff. | | −0.58*** | −0.65*** | −0.84*** | −0.88*** | −1.25*** | −1.85*** | −2.39*** | −2.94*** | −4.09*** | −5.39*** | −6.96*** | −9.12*** |
| | | | | | | | $PNBOC >$ 1std | | | | | | |
| Avg. Ret. | 10.09*** | 12.82*** | 15.78*** | 18.67*** | 21.59*** | 24.84*** | 28.31*** | 32.09*** | 36.39*** | 40.93*** | 45.77*** | 51.11*** | 55.98*** |
| Avg. Diff. | | 2.73*** | 2.96*** | 2.89*** | 2.92*** | 3.25*** | 3.46*** | 3.78*** | 4.30*** | 4.54*** | 4.84*** | 5.35*** | 4.86** |
| | | | | | | | $PNBOC < -$1std | | | | | | |
| Avg. Ret. | −4.48* | −6.35** | −8.27*** | −10.42*** | −12.57*** | −14.86*** | −17.97*** | −21.69*** | −25.49*** | −29.75*** | −34.27*** | −39.40*** | −45.73*** |
| Avg. Diff. | | −1.87*** | −1.92*** | −2.14*** | −2.15*** | −2.29*** | −3.11*** | −3.72*** | −3.80*** | −4.26*** | −4.51*** | −5.14*** | −6.32*** |

This table reports the average 30-day call option returns (in %) for all months in the sample and conditioned to months where $PNBOC$ is positive, negative, above one standard deviation of its time series and below minus one standard deviation. Also reported is the average difference (in %) between the returns of the current column and the previous one. Call options have deltas ranging from 0.2 to 0.8. The number of observations for each of the panels, from top to bottom, is: 6293, 3003, 3290, 753 and 613. *, ** and *** represent statistical significance at the 10%, 5% and 1% levels, respectively. Statistical significance is based on 25,000 bootstrap draws from the respective sample of option returns (and pairwise bootstrap draws for the return differences). The sample ranges from January 4, 1996 to December 31, 2020.

Table D.2: **Option net demand and average put option returns**

| Delta | −0.20 | −0.25 | −0.30 | −0.35 | −0.40 | −0.45 | −0.50 | −0.55 | −0.60 | −0.65 | −0.70 | −0.75 | −0.80 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | Whole Sample | | | | | | |
| Avg. Ret. | −53.21*** | −45.94*** | −40.49*** | −36.34*** | −32.78*** | −29.92*** | −27.61*** | −25.50*** | −23.52*** | −21.68*** | −19.89*** | −18.14*** | −16.31*** |
| Avg. Diff. | | 7.27*** | 5.45*** | 4.16*** | 3.56*** | 2.86*** | 2.31*** | 2.11*** | 1.98*** | 1.83*** | 1.79*** | 1.75*** | 1.83*** |
| | | | | | | | $PNBOP > 0$ | | | | | | |
| Avg. Ret. | −39.42*** | −33.10*** | −28.45*** | −25.13*** | −22.40*** | −20.29*** | −18.72*** | −17.35*** | −16.10*** | −14.90*** | −13.68*** | −12.40*** | −11.03*** |
| Avg. Diff. | | 6.31*** | 4.65*** | 3.32*** | 2.73*** | 2.12*** | 1.57*** | 1.37*** | 1.25*** | 1.19*** | 1.22*** | 1.28*** | 1.38*** |
| | | | | | | | $PNBOP < 0$ | | | | | | |
| Avg. Ret. | −68.52*** | −60.18*** | −53.85*** | −48.77*** | −44.29*** | −40.61*** | −37.47*** | −34.55*** | −31.75*** | −29.21*** | −26.78*** | −24.51*** | −22.17*** |
| Avg. Diff. | | 8.34*** | 6.33*** | 5.08*** | 4.48*** | 3.68*** | 3.14*** | 2.92*** | 2.80*** | 2.54*** | 2.43*** | 2.27*** | 2.34*** |
| | | | | | | | $PNBOP >$ 1std | | | | | | |
| Avg. Ret. | 60.45*** | 53.14*** | 46.08*** | 40.06*** | 34.82*** | 30.11*** | 25.67** | 21.54** | 17.72** | 14.08* | 10.84 | 7.96 | 5.48 |
| Avg. Diff. | | −7.31** | −7.06*** | −6.01*** | −5.24*** | −4.71*** | −4.44*** | −4.13*** | −3.82*** | −3.65*** | −3.24*** | −2.88*** | −2.48*** |
| | | | | | | | $PNBOP <$ −1std | | | | | | |
| Avg. Ret. | −76.62*** | −68.81*** | −62.76*** | −57.23*** | −52.23*** | −48.19*** | −45.00*** | −42.36*** | −39.63*** | −36.95*** | −34.32*** | −31.92*** | −29.37*** |
| Avg. Diff. | | 7.81*** | 6.05*** | 5.52*** | 5.00*** | 4.04*** | 3.19*** | 2.64*** | 2.73*** | 2.68*** | 2.64*** | 2.39*** | 2.55*** |

This table reports the average 30-day put option returns (in %) for all months in the sample and conditioned to months where $PNBOP$ is positive, negative, above one standard deviation of its time series and below minus one standard deviation. Also reported is the average difference (in %) between the returns of the current column and the previous one. Put options have deltas ranging from -0.8 to -0.2. The number of observations for each of the panels, from top to bottom, is: 6293, 3310, 2983, 629 and 844. *, ** and *** represent statistical significance at the 10%, 5% and 1% levels, respectively. Statistical significance is based on 25,000 bootstrap draws from the respective sample of option returns (and pairwise bootstrap draws for the return differences). The sample ranges from January 4, 1996 to December 31, 2020.

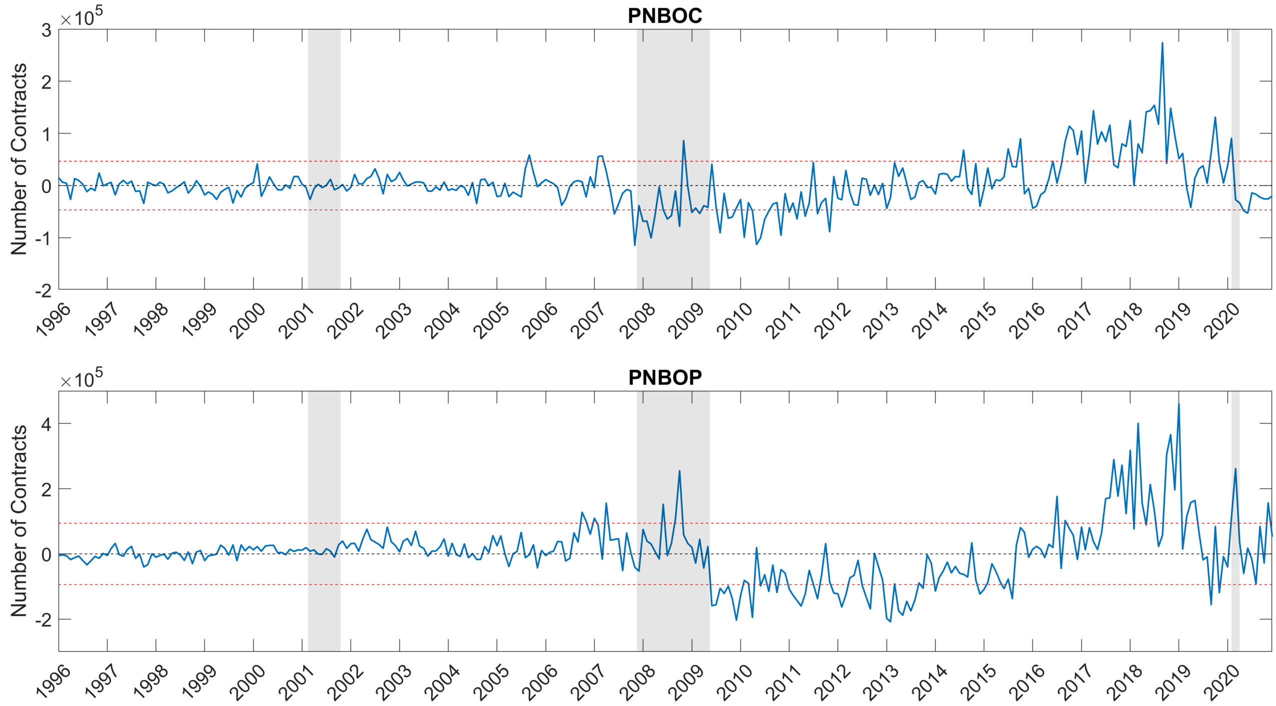


Fig. 1: **PNBOC and PNBOP.** This figure plots in the upper (lower) panel the monthly time series of $PNBOC$ ($PNBOP$). The red dashed lines denote one standard deviation and minus one standard deviation of $PNBOC$ ($PNBOP$). Shaded areas depict National Bureau of Economic Research (NBER) recession dates. The sample ranges from January, 1996 to December, 2020.

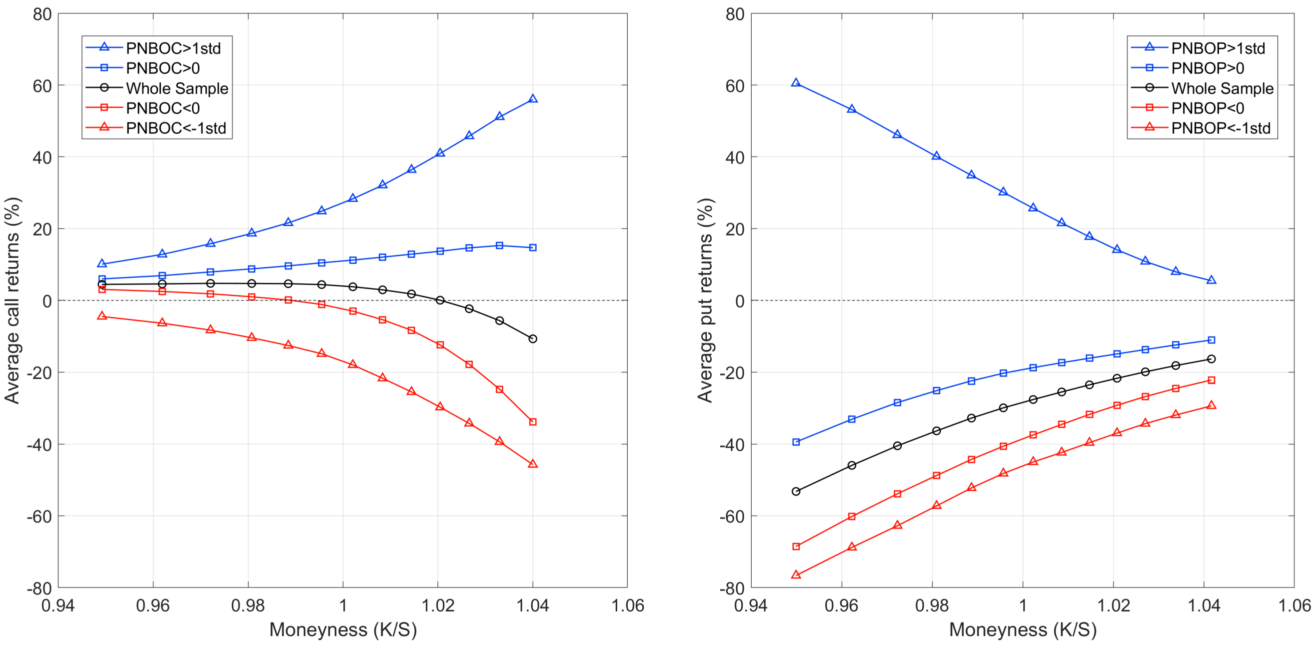


Fig. 2: **Average option returns and PNBO.** This figure plots in the left (right) panel average 30-day call (put) returns calculated over the whole sample and conditioned to months where $PNBOC$ ($PNBOP$) is: above one standard deviation of its time series, above zero, below zero and below minus one standard deviation. Call (put) options have deltas ranging from 0.2 to 0.8 (-0.8 to -0.2) and their average returns are plotted against the corresponding average moneyness ($K/S$) over the whole sample. The sample ranges from January 4, 1996 to December 31, 2020.

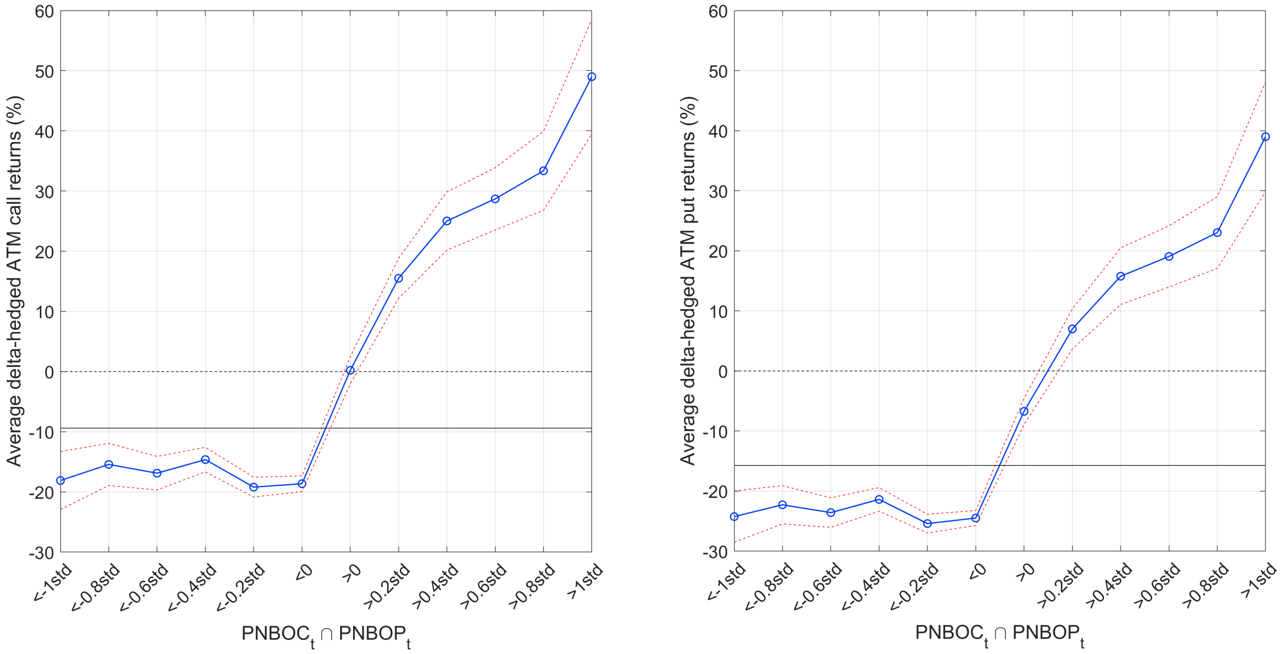


Fig. 3: **Average delta-hedged option returns and PNBO.** This figure plots in the left (right) panel average 30-day delta-hedged ATM call (put) returns calculated conditioned to months where both $PNBOC_t$ and $PNBOP_t$ are: below several thresholds (ranging from minus one standard deviation of their time series to zero) and above several thresholds (ranging from zero to one standard deviation of their time series). The red dashed lines denote one-standard error bands based on 1,000 bootstrap draws from the respective sample of returns. The black horizontal line is the average return over the whole sample. The sample ranges from January, 1996 to December, 2020.

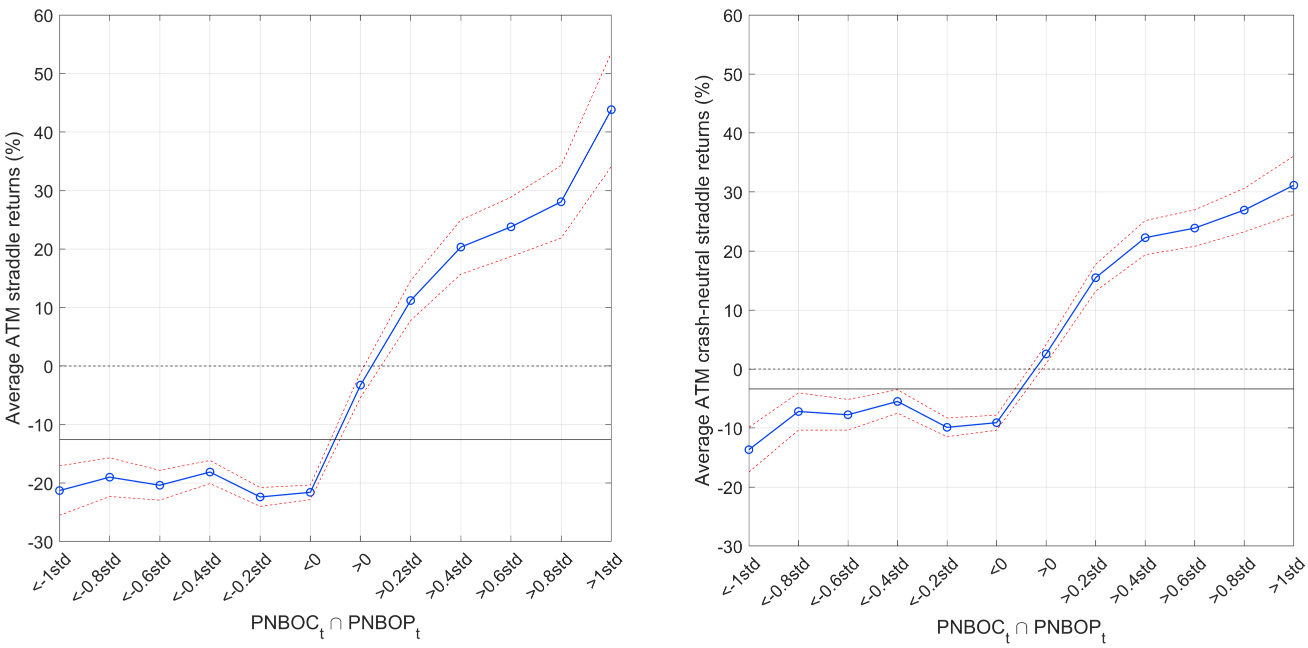


Fig. 4: **Average straddle returns and PNBO.** This figure plots in the left (right) panel average 30-day straddle (crash-neutral straddle) returns calculated conditioned to months where both $PNBOC_t$ and $PNBOP_t$ are: below several thresholds (ranging from minus one standard deviation of their time series to zero) and above several thresholds (ranging from zero to one standard deviation of their time series). The red dashed lines denote one-standard error bands based on 1,000 bootstrap draws from the respective sample of returns. The black horizontal line is the average return over the whole sample. The sample ranges from January, 1996 to December, 2020.

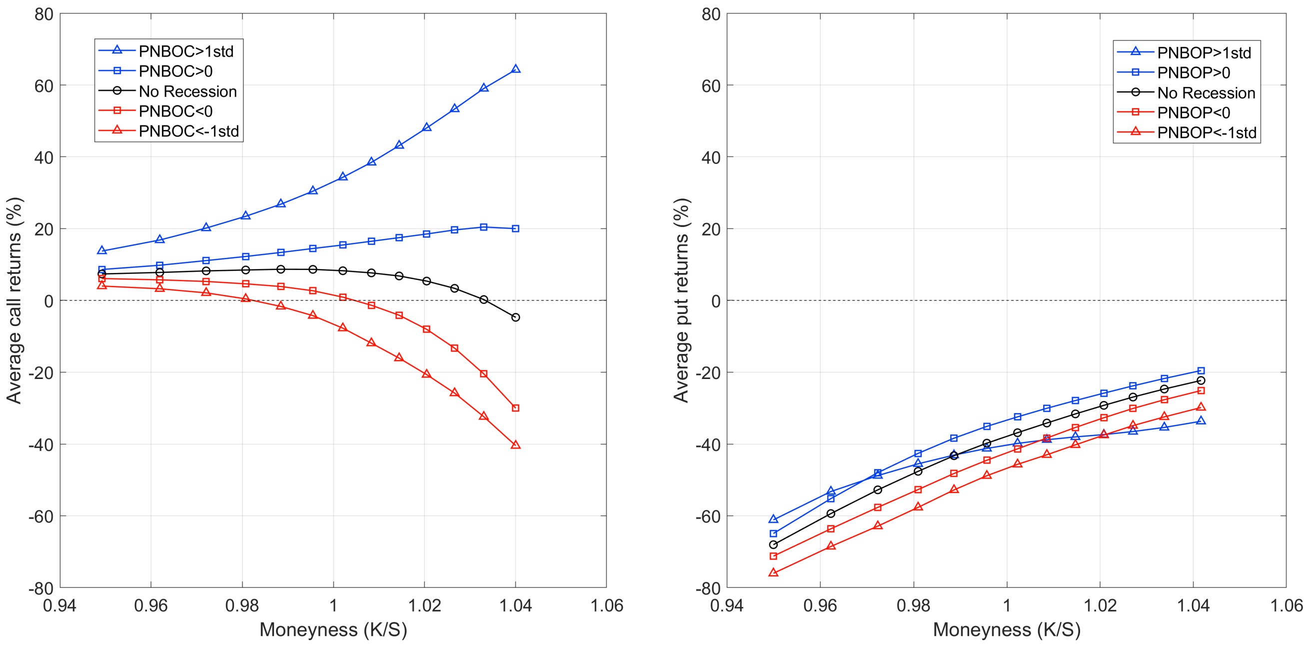


Fig. 5: **Average option returns and PNBO controlling for no recession.** This figure plots in the left (right) panel average 30-day call (put) returns calculated over months without recession and conditioned to months where there is no recession and $PNBOC$ ($PNBOP$) is: above one standard deviation of its time series in this sample, above zero, below zero and below minus one standard deviation. Call (put) options have deltas ranging from 0.2 to 0.8 (-0.8 to -0.2) and their average returns are plotted against the corresponding average moneyness ($K/S$) over the whole sample. The sample ranges from January 4, 1996 to December 31, 2020.

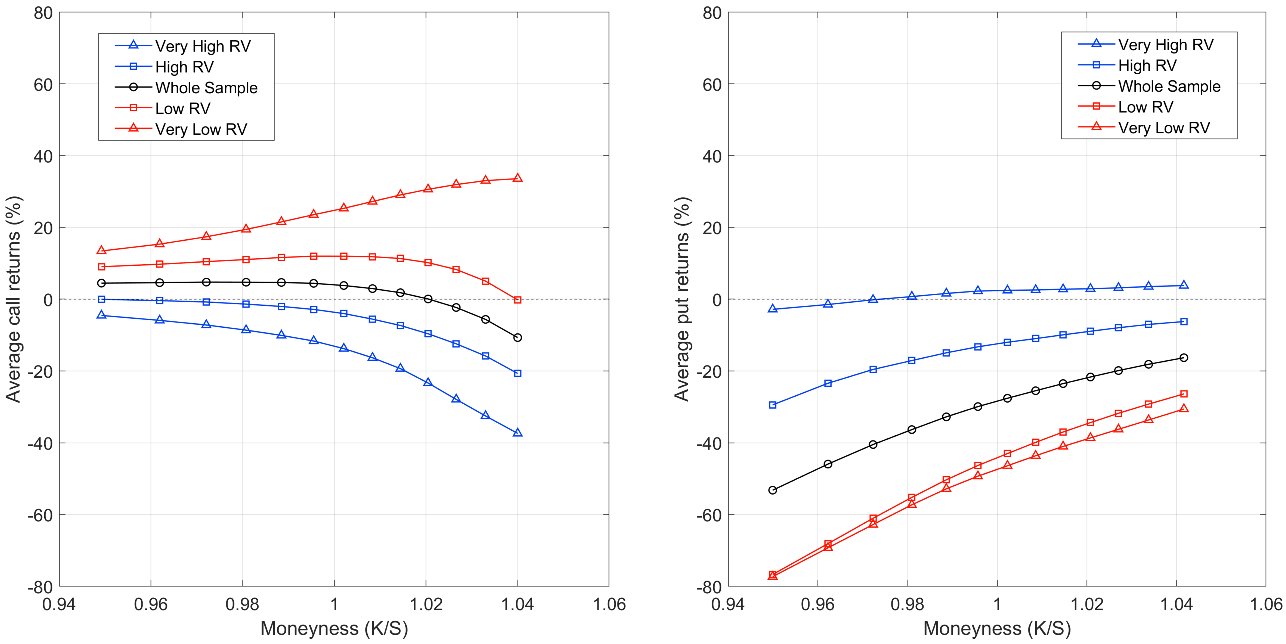


Fig. 6: **Average option returns and $RV$.** This figure plots in the left (right) panel average 30-day call (put) returns calculated over the whole sample and conditioned to months where $RV$ is: above its 80% quantile (very high), above its median (high), below its median (low) and below its 20% quantile (very low). Call (put) options have deltas ranging from 0.2 to 0.8 (-0.8 to -0.2) and their average returns are plotted against the corresponding average moneyness ($K/S$) over the whole sample. The sample ranges from January 4, 1996 to December 31, 2020.

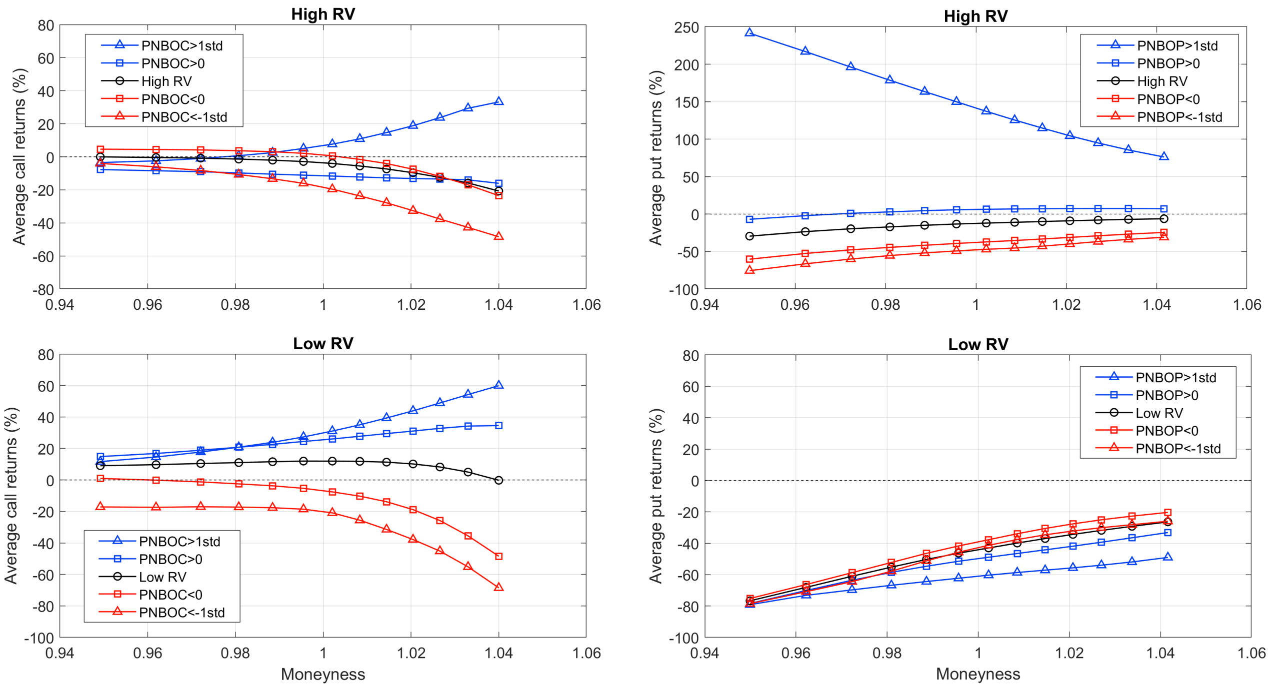


Fig. 7: **Average option returns and PNBO controlling for $RV$.** This figure plots in the upper left (right) panel average 30-day call (put) returns calculated over months where $RV$ is above its median (high) and conditioned to months where $RV$ is above its median and $PNBOC$ ($PNBOP$) is: above one standard deviation of its time series in this sample, above zero, below zero and below minus one standard deviation. The lower panels plot analogous results when $RV$ is below its median (low). Call (put) options have deltas ranging from 0.2 to 0.8 (-0.8 to -0.2) and their average returns are plotted against the corresponding average moneyness ($K/S$) over the whole sample. The sample ranges from January 4, 1996 to December 31, 2020.

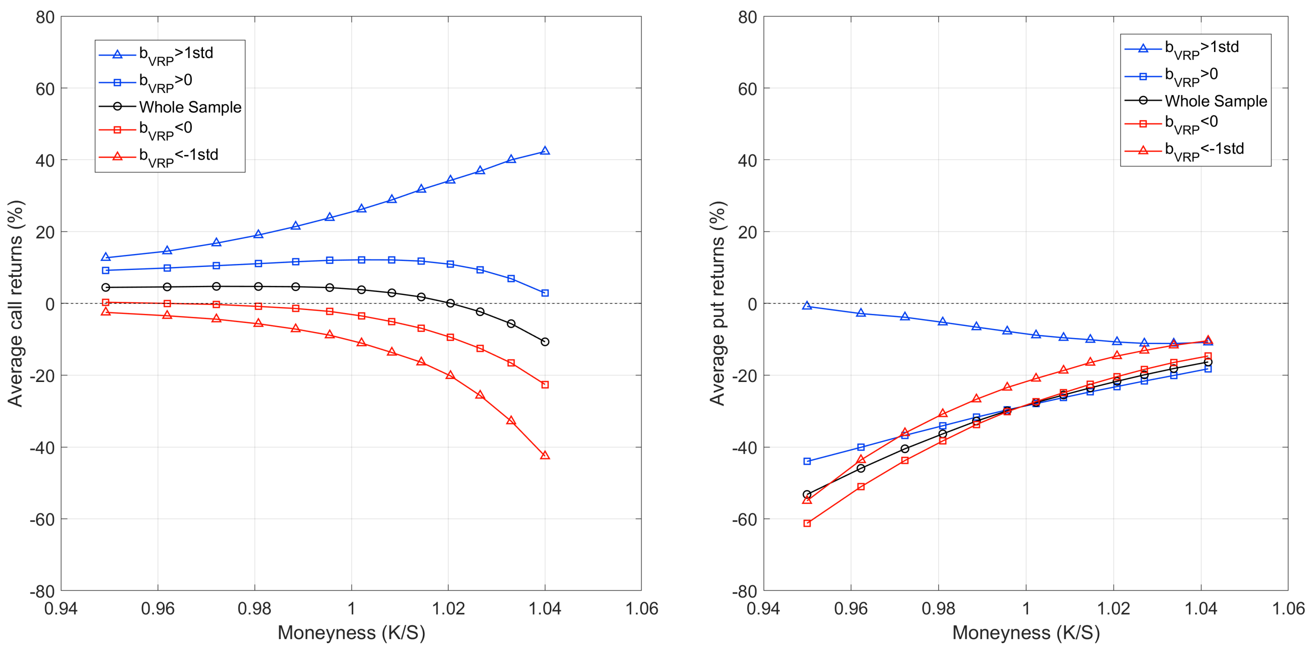


Fig. 8: **Average option returns and $b_{VRP}$.** This figure plots in the left (right) panel average 30-day call (put) returns calculated over the whole sample and conditioned to months where $b_{VRP}$ is: above one standard deviation of its time series, above zero, below zero and below minus one standard deviation. Call (put) options have deltas ranging from 0.2 to 0.8 (-0.8 to -0.2) and their average returns are plotted against the corresponding average moneyness ($K/S$) over the whole sample. The sample ranges from January 4, 1996 to December 31, 2020.

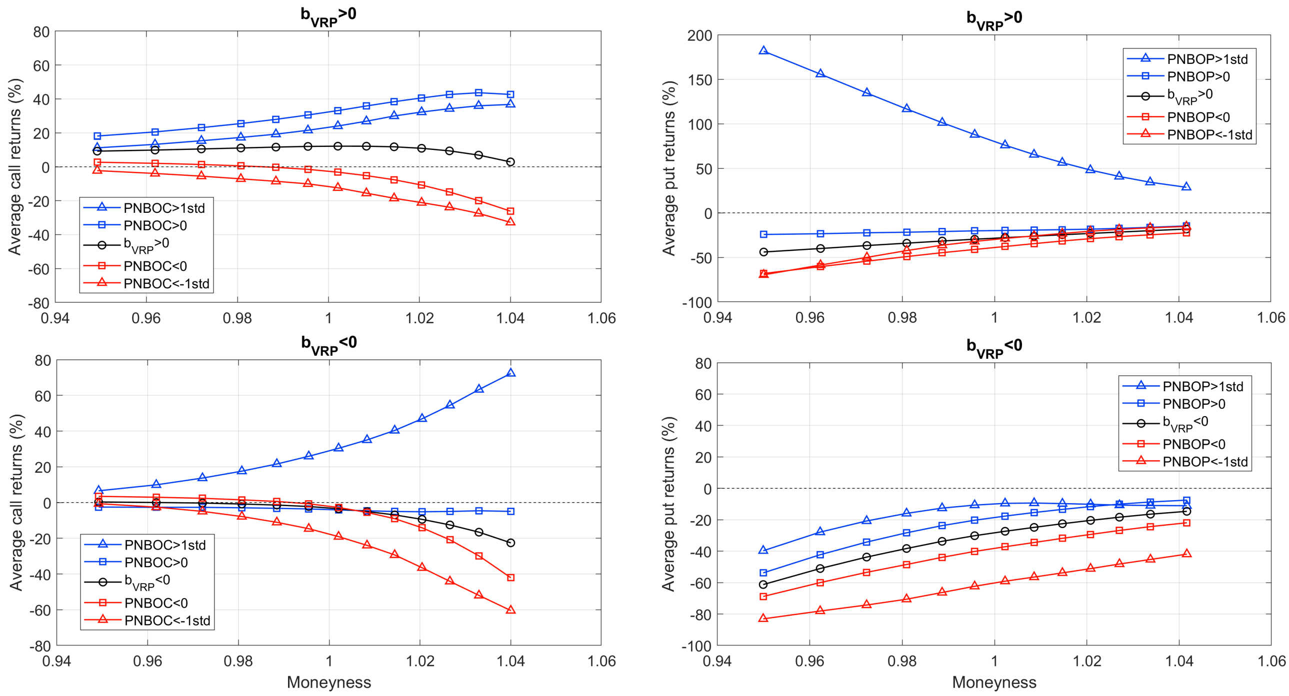


Fig. 9: **Average option returns and PNBO controlling for $b_{VRP}$.** This figure plots in the upper left (right) panel average 30-day call (put) returns calculated over months where $b_{VRP}$ is above zero and conditioned to months where $b_{VRP}$ is above zero and $PNBOC$ ($PNBOP$) is: above one standard deviation of its time series in this sample, above zero, below zero and below minus one standard deviation. The lower panels plot analogous results when $b_{VRP}$ is below zero. Call (put) options have deltas ranging from 0.2 to 0.8 (-0.8 to -0.2) and their average returns are plotted against the corresponding average moneyness ($K/S$) over the whole sample. The sample ranges from January 4, 1996 to December 31, 2020.

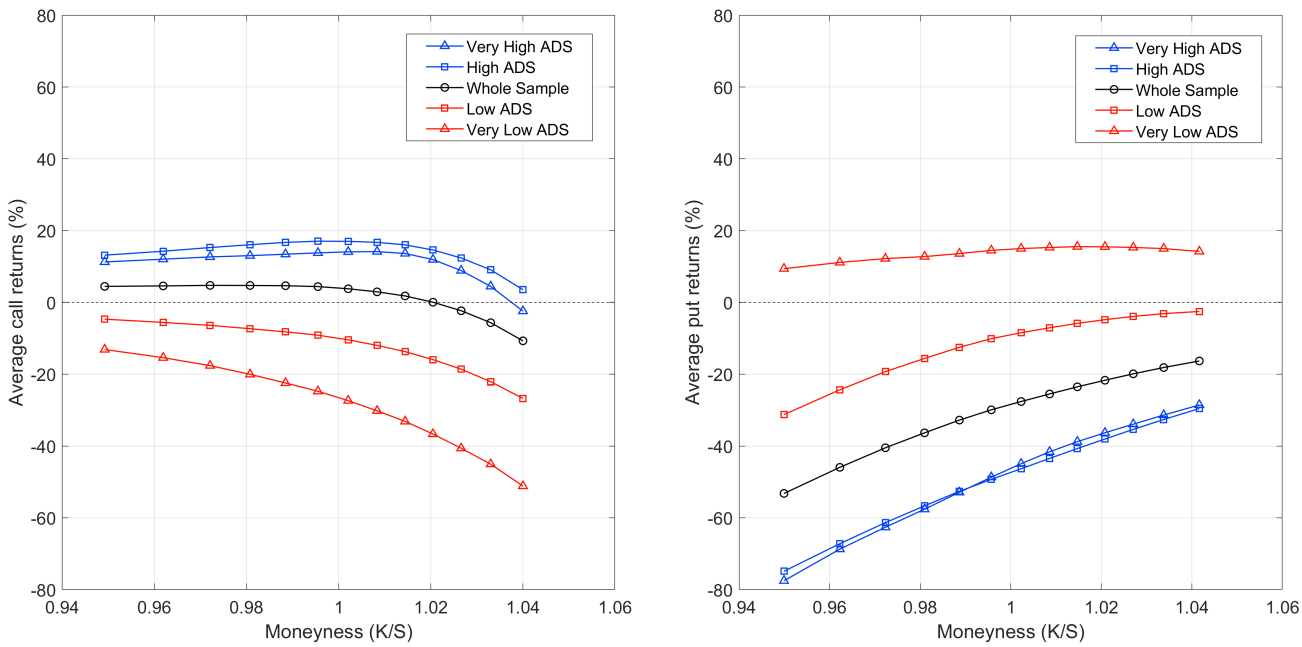


Fig. 10: **Average option returns and** $\boldsymbol{ADS}$**.** This figure plots in the left (right) panel average 30-day call (put) returns calculated over the whole sample and conditioned to months where $ADS$ is: above its 80% quantile (very high), above its median (high), below its median (low) and below its 20% quantile (very low). Call (put) options have deltas ranging from 0.2 to 0.8 (-0.8 to -0.2) and their average returns are plotted against the corresponding average moneyness $(K/S)$ over the whole sample. The sample ranges from January 4, 1996 to December 31, 2020.

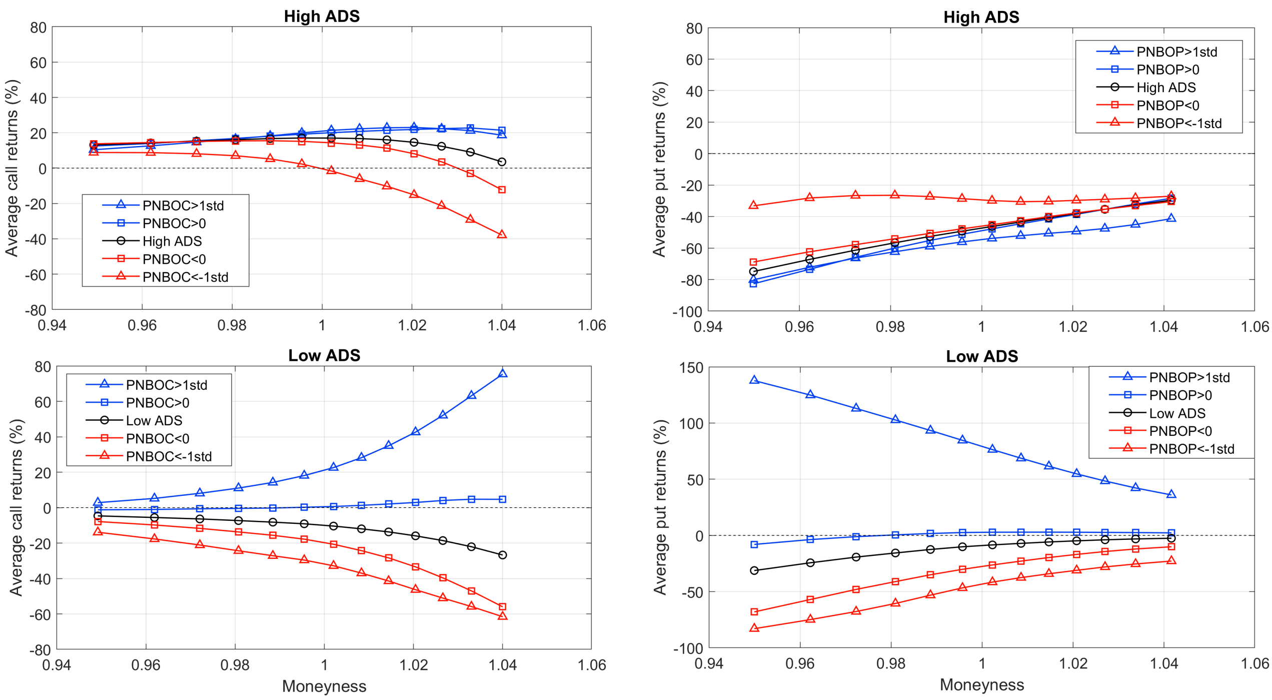


Fig. 11: **Average option returns and PNBO controlling for** $ADS$**.** This figure plots in the upper left (right) panel average 30-day call (put) returns calculated over months where $ADS$ is above its median (high) and conditioned to months where $ADS$ is above its median and $PNBOC$ ($PNBOP$) is: above one standard deviation of its time series in this sample, above zero, below zero and below minus one standard deviation. The lower panels plot analogous results when $ADS$ is below its median (low). Call (put) options have deltas ranging from 0.2 to 0.8 (-0.8 to -0.2) and their average returns are plotted against the corresponding average moneyness ($K/S$) over the whole sample. The sample ranges from January 4, 1996 to December 31, 2020.

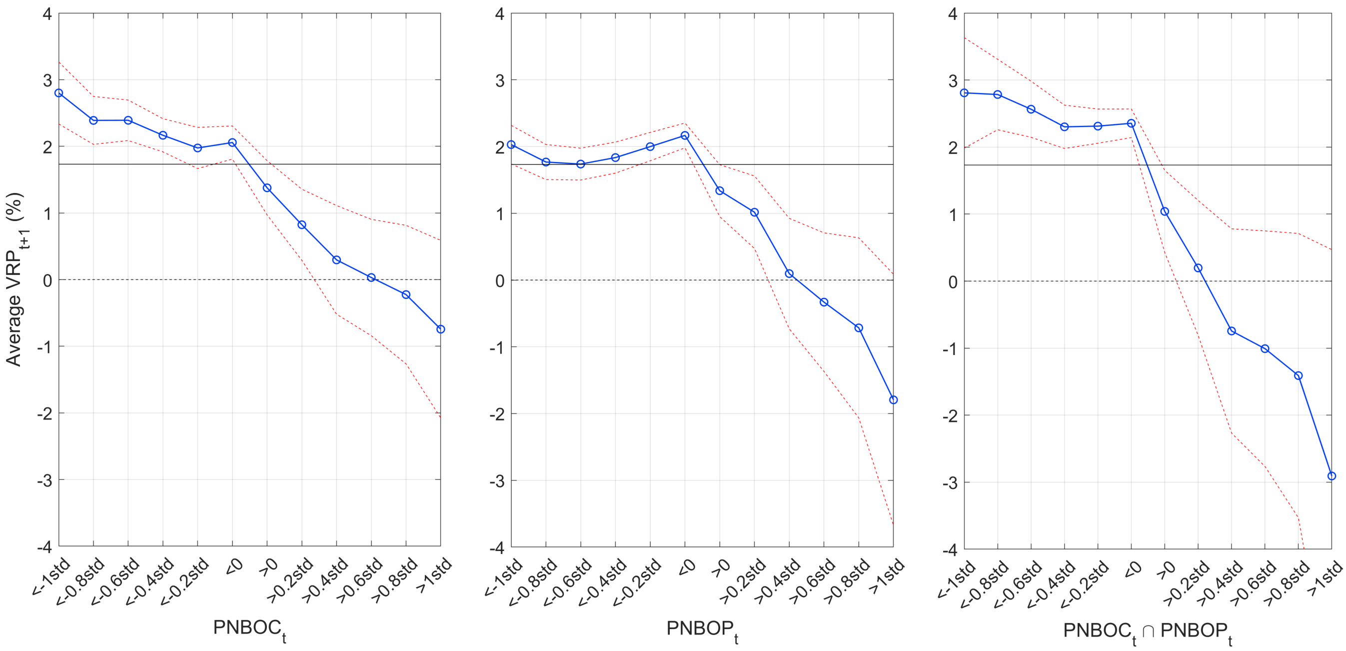


Fig. 12: **Variance risk premium and PNBO.** This figure plots the average of the variance risk premium calculated in the months $t+1$ following the months $t$ where $PNBOC_t$ is, $PNBOP_t$ is, and both $PNBOC_t$ and $PNBOP_t$ are, respectively: below several thresholds (ranging from minus one standard deviation to zero) and above several thresholds (ranging from zero to one standard deviation). The red dashed lines denote one-standard error bands based on 1,000 bootstrap draws from the respective sample of $VRP_{t+1}$. The black horizontal line is the average $VRP_{t+1}$ over the whole sample. The sample ranges from January, 1996 to December, 2020.

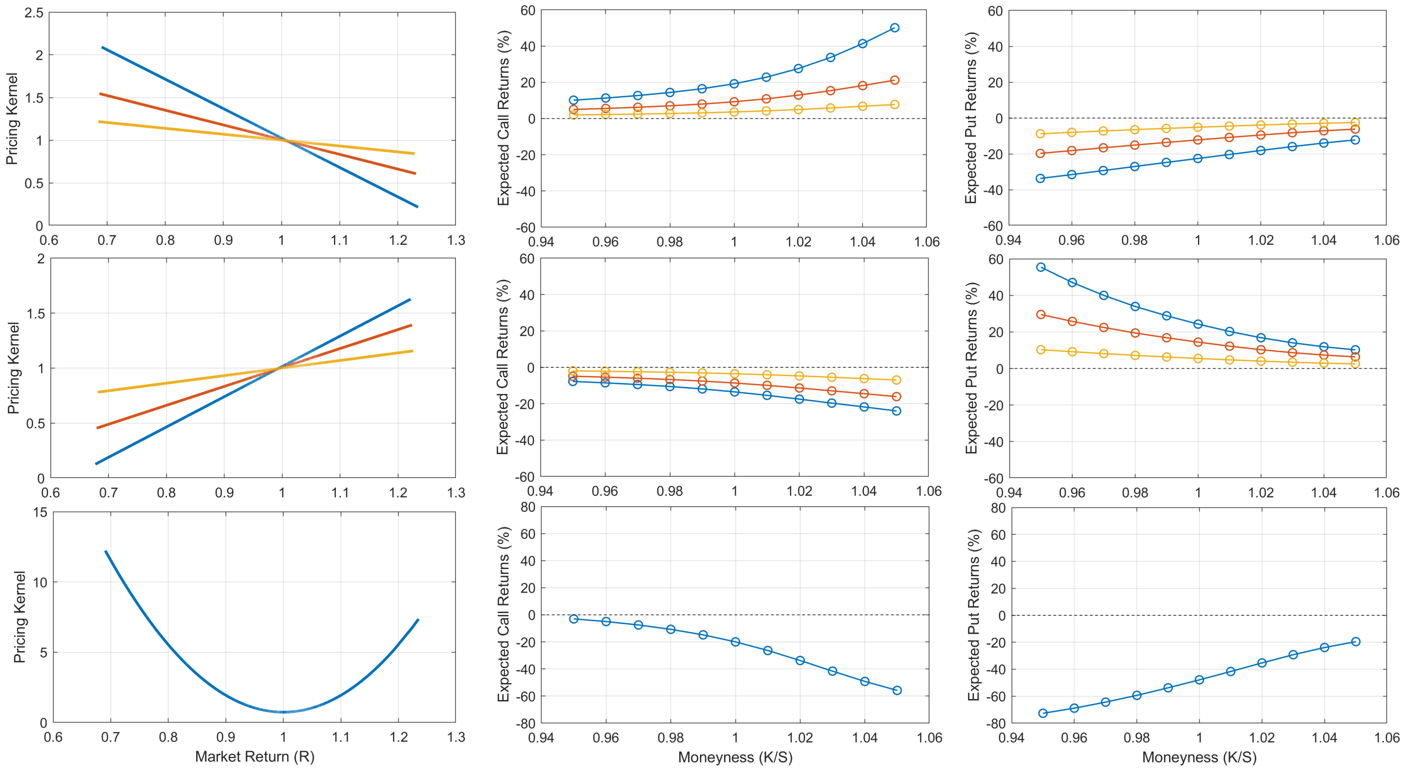


Fig. C.1: **Hypothetical pricing kernels and expected option returns.** This figure plots hypothetical pricing kernels as a function of market returns in the left panel, and the implied expected call (put) returns in the middle (right) panel. To construct the figure, we rely on a daily sample of overlapping 30-calendar day market returns between January 4, 1996 and December 31, 2020 (i.e., we take the unconditional physical distribution). We specify different pricing kernels as a function of those market returns (upper panel: linear and monotonically decreasing, with different slopes; middle panel: linear and monotonically increasing, with different slopes; lower panel: U-shaped). Then, we report the corresponding implied expected call and put returns (with the same color), which are computed as the ratio between the expected option payoff under the physical measure and the expected option payoff under the risk-neutral measure (implied by the pricing kernel).