Journal of Empirical Finance 69 (2022) 285–302

Contents lists available at ScienceDirect

Journal of Empirical Finance

journal homepage: www.elsevier.com/locate/jempfin

Consumption risks in option returns[☆]

Shuwen Yang [c,*], Kevin Aretz [a], Hening Liu [a,b], Yuzhao Zhang [d]

[a] Accounting and Finance Group, Alliance Manchester Business School, the University of Manchester, Booth Street West, Manchester M15 6PB, UK

[b] School of Economics and Management, China University of Mining and Technology, No. 1, Daxue Road, Xuzhou, 221116, PR China

[c] PBC School of Finance, Tsinghua University, 43 Chengfu Road, Haidian District, Beijing, PR China

[d] AllianceBernstein L.P., New York, NY, 10105, USA

ARTICLE INFO

JEL classification:
G11
G12

Keywords:
Consumption growth
Option returns
Recursive utility
Volatility risk

ABSTRACT

We offer evidence that exposures to consumption growth and consumption volatility are significantly priced in the cross-section of delta-hedged option returns. Consumption growth commands a positive risk premium, whereas consumption volatility commands a negative risk premium. Our results suggest that consumption risk exposures provide rational foundations for well-known relations between options moneyness or idiosyncratic underlying-stock volatility and the cross-section of delta-hedged option returns. Furthermore, those risk premiums can also price stocks. In a representative-agent economy with recursive preferences, our results suggest that investors prefer early resolution of uncertainty.

1. Introduction

In the standard consumption-based asset pricing model (CCAPM) pioneered by Breeden (1979), the risk premium on an asset is a multiple of its exposure to consumption risk, the covariance of the asset return with contemporaneous aggregate consumption growth. In long-run risk models with Epstein and Zin's (1989) recursive preferences and richer dynamics of consumption growth (e.g., Bansal and Yaron (2004) and Gallant et al. (2019)), both expected consumption growth and consumption volatility are also priced. Most of the early studies in the consumption-based asset pricing literature focus on the impact of the first moment of consumption growth on stocks (e.g., Lettau and Ludvigson (2001), Parker and Julliard (2005), and Yogo (2006)). More recently, Boguth and Kuehn (2013) stress the importance of consumption volatility for stocks. However, the consumption-based framework is, in theory, applicable to all traded assets, including options. Delta-hedged options are particularly sensitive to the underlying asset's volatility, which is in turn determined by the fundamental consumption volatility. Option returns are related to consumption volatility in light of two strands of literature. First, empirical studies (Coval and Shumway (2001), Goyal and Saretto (2009), Cao and Han (2013), and Hu and Jacobs (2020)) show that stock volatility risk is priced in option returns. Second, theoretical analysis (Veronesi (1999), Bansal and Yaron (2004), and Drechsler and Yaron (2011)) suggests that consumption volatility is

☆ We thank Rossen Valkanov (the editor), an anonymous associate editor, an anonymous referee, Michael Brennan, Sungjun Cho, Amit Goyal, Kai Li and seminar participants at Shanghai University of Finance and Economics, the University of Nottingham Ningbo China, the 2018 Frontiers of Factor Investing Conference (Lancaster), the 2018 China Meeting of the Econometric Society (Shanghai), the 2018 Lingnan Macro Workshop on Financial Risks and Macroprudencial Policy (Guangzhou), the 2018 Guanghua International Symposium on Finance (Beijing), the 2018 Econometric Conference on Big Data and AI in Honor of Ronald Gallant (Beijing), the 2019 International Accounting and Finance Doctoral Symposium (Milano), the 2019 European Meeting of the Econometric Society (Manchester), and the 2019 China Finance Review International Conference (Shanghai) for helpful comments and suggestions. The views expressed in this paper are those of the authors and do not necessarily reflect those of AllianceBernstein.

* Corresponding author.
E-mail addresses: yangshw@pbcsf.tsinghua.edu.cn (S. Yang), kevin.aretz@manchester.ac.uk (K. Aretz), hening.liu@manchester.ac.uk (H. Liu), yuzhao.zhang@gmail.com (Y. Zhang).

https://doi.org/10.1016/j.jempfin.2022.10.001
Received 7 August 2020; Received in revised form 20 July 2022; Accepted 1 October 2022
Available online 7 October 2022
0927-5398/

important for explaining the behavior of stock volatility. The strong connection between options and volatility provides us with powerful test assets to identify the consumption volatility premium.

In our paper, we evaluate the ability of consumption risks to price delta-hedged call options. We show that consumption risks offer rational explanations for the following well-known option anomalies. Bakshi and Kapadia (2003) document that delta-hedged call returns are negative and low moneyness (out-of-the-money (OTM)) options have more negative returns. Cao and Han (2013) discover that delta-hedged option returns are negative and decreasing in the idiosyncratic volatility of the underlying stock (IVOL). The option portfolio with low moneyness and high IVOL loads more negatively on consumption growth and more positively on consumption volatility than the option portfolio with high moneyness (in-the-money (ITM)) and low IVOL does. Using options as test assets, we identify highly significant consumption growth and consumption volatility risk premiums which can also price the cross-section of stock returns sorted on well-known characteristics including the market beta, size and valuation ratios. The significant consumption growth premium supports the CCAPM. The negative and significant risk premium on consumption volatility supports the recursive utility model featuring a preference for early resolution of uncertainty. That is, the elasticity of intertemporal substitution (EIS) of the representative agent exceeds the inverse of relative risk aversion (RRA).

To better understand how consumption risks affect options, we study a delta-hedged call option in the representative-agent model of Lettau et al. (2008). In that model, consumption growth follows a Markov-switching process in which the mean growth rate and the volatility of the innovation shock are characterized by two independent Markov chains. The representative agent has Epstein and Zin's (1989) recursive preferences that disentangle RRA and EIS. The agent cannot observe the state of the economy but has to learn about it by observing realized consumption growth. The agent updates beliefs according to Bayes' rule. It can be shown that the log-linearized pricing kernel is an affine function of consumption growth and its conditional mean and conditional volatility. We adopt this consumption process instead of the long-run risk model in Bansal and Yaron (2004) because Bayesian learning embedded in the model generates time-series estimates of expected consumption growth and consumption volatility. Using the impulse response analysis, we show that when the EIS is greater than the inverse of RRA, a realized shock to consumption growth lowers the price of a call option written on the underlying stock despite raising the option's implied volatility. The call option's loss, however, is lower than the gain from shorting delta stocks because the call option price is convex in the underlying stock's price and because the implied volatility increases in response to the shock. Thus, the value of the delta-hedged option (which is long the call option and short delta stocks) increases, implying that the delta-hedged option is negatively exposed to consumption growth and its conditional mean but is positively exposed to consumption volatility.

To test these asset pricing implications, we estimate the Markov-switching model and obtain estimates of the conditional mean and volatility of consumption growth. We sort options into portfolios based on stock or option characteristics that generate a spread in average option returns and test whether the spread in returns can be explained by betas with respect to sources of consumption risk. We use firm-level IVOL and option moneyness as sorting variables because Cao and Han (2013) show that IVOL is negatively related to delta-hedged option returns, and in addition, Bakshi and Kapadia (2003) discover a positive relation between delta-hedged call option returns and options moneyness. To match the quarterly consumption data, we choose a cross-section of options with times-to-maturity between about three to six months at the end of each quarter. Our time-to-maturity choice guarantees that the options expire after the end of the coming quarter and confines the times-to-maturity to be within a reasonable range. We then compute the quarterly return of a portfolio that buys one call option and delta-hedges it with the underlying stock. Delta-hedging the option neutralizes the effect of movements in the underlying stock's price on option returns, ensuring that our results do not simply hinge on stock returns. We finally form 16 equally-weighted delta-hedged call option portfolios sorted independently on the underlying stock's IVOL and options moneyness.

In line with existing studies (e.g., Cao and Han (2013)), the mean returns of our constructed delta-hedged option portfolios are all negative. We provide a macroeconomic-based explanation for this result. Most of the portfolios have negative exposures toward consumption growth. According to the CCAPM, investors accept a lower or even negative return when an asset pays off in adverse macroeconomic conditions. Moreover, we find that our constructed delta-hedged option portfolios all have positive consumption volatility exposures. Taken together, our results suggest that the delta-hedged option portfolios are countercyclical assets accommodating investors' hedging concerns and thus have negative returns on average. We also find that the exposures of the option portfolios to consumption growth risk (consumption volatility risk) become more negative (positive) at higher IVOL levels, implying that options written on higher IVOL stocks offer better protection against adverse conditions featuring low consumption growth or high uncertainty.

Using both Fama–MacBeth (Fama and MacBeth (1973), henceforth FM) regressions and Hansen's (1982) generalized method of moments (GMM), we find that the consumption growth and volatility risk exposures are priced in options and the two methods generate consistent risk premium estimates. The estimated consumption growth risk premium is positive and significant, while the estimated consumption volatility risk premium is negative and significant. The intercepts from the FM regressions are all insignificantly different from zero. The GMM overidentifying restrictions test fails to reject the Euler equation, which is based on the observation that the stochastic discount factor (SDF) is approximately affine in the state variables. These results are consistent with Jagannathan and Wang (2007) and Boguth and Kuehn (2013) who investigate the cross-section of stock returns. The two sources of consumption risks both contribute to the negative delta-hedged option returns and explain over 45% of the cross-sectional variation in option returns. In relation to the consumption-based model, the positive premium on consumption growth suggests that investors are risk averse, while the negative premium on consumption volatility implies that investors prefer early resolution of uncertainty (EIS > 1/RRA).

To examine whether consumption risk premiums estimated from options can price stocks, we first compare the average returns of the market portfolio and the Fama–French 25 size-value stock portfolios with those predicted by our estimated consumption risk premiums. We estimate each stock portfolio's exposures to consumption growth and the changes in mean growth and consumption volatility, and calculate the predicted mean returns of each portfolio via combining those exposures with the consumption risk premiums estimated from options. The correlation between the mean portfolio returns in the data and those implied by our estimates of the consumption risk premiums is 0.54. More formally, we use GMM to test implications from the Euler equation using the joint cross-section of the 16 option portfolios and the 25 size-value stock portfolios as test assets. The consumption growth and its volatility are significantly priced, with risk premiums being quantitatively similar to those obtained using only option portfolios as test assets. Moreover, the expected consumption growth is also significantly priced, which indicates that adding stock portfolios helps us to identify prices of shocks to expected consumption growth.

Our work contributes to the literature on the cross-section of option returns as well as studies on time-varying economic uncertainty. Cao and Han (2013) show that delta-hedged option returns decrease with the underlying stock's IVOL. They argue that options written on high IVOL stocks are more difficult to hedge, inducing dealers to charge a higher premium in the presence of limits to arbitrage. The question remains why investors are willing to pay the extra premium. Our analysis complements theirs by showing that the options written on high IVOL stocks provide a better hedge against adverse macroeconomic conditions, making investors willing to accept low or even negative returns. Hu and Jacobs (2020) show that returns on call (put) stock-option portfolios decrease (increase) with the underlying stock volatility.[1] By mainly using variance swaps, Dew-Becker et al. (2017) find that shocks to expected aggregate stock market volatility are not priced beyond one quarter. Our work is different from theirs in the following aspects: (1) Dew-Becker et al. (2017) study variance swaps and synthetic variance swaps on the S&P 500 index, while we study the cross-section of options written on individual stocks; (2) we use consumption data to estimate the time-series of consumption volatility and the associated risk premium rather than assessing the performance of calibrated consumption-based models in matching the term structure of variance risk implied by variance swaps. Existing studies such as Drechsler and Yaron (2011), Boguth and Kuehn (2013), Romeo (2015), and Liu and Zhang (2022) emphasize the importance of time-varying economic uncertainty and a preference for early resolution of uncertainty in explaining the behavior of stock returns. Our paper complements these studies by providing empirical support using option returns.

The rest of our paper is organized as follows. Section 2 introduces the theoretical model motivating our empirical analysis. In Section 3, we use numerical analysis to study the impacts of consumption risks on delta-hedged option returns. In Section 4, we test whether loadings on consumption growth and changes in its conditional moments forecast the cross-section of delta-hedged option returns. Section 5 summarizes and concludes. The Internet Appendix contains additional derivations and empirical results.

2. The model

In this section, we briefly introduce the consumption-based model of Lettau et al. (2008). In the model, the growth rate of consumption follows a Markov-switching process, and the representative agent has the recursive preferences of Epstein and Zin (1989). We next follow Boguth and Kuehn (2013) in linearizing the SDF to derive an equation for expected returns. We use the model to guide our empirical analysis.

2.1. Consumption dynamics

We follow McConnell and Perez-Quiros (2000) and Lettau et al. (2008) and assume that consumption growth follows a Markov-switching process in which the conditional mean and volatility states follow two independent Markov chains. More specifically, we assume that the log consumption growth rate, Δc_{t+1}, follows the process

$$\Delta c_{t+1} \equiv \ln\left(\frac{C_{t+1}}{C_t}\right) = \mu_t + \sigma_t \epsilon_{t+1}, \quad \epsilon_{t+1} \sim N(0,1), \tag{1}$$

where C_t is consumption at time t, and μ_t the conditional mean and σ_t the conditional volatility of Δc_{t+1}. We assume two states for mean growth, $\mu_t \in \{\mu_l, \mu_h\}$, and two states for the volatility of the innovation shock, $\sigma_t \in \{\sigma_l, \sigma_h\}$. The transition matrix for the mean and volatility states are $\boldsymbol{P}^\mu$ and $\boldsymbol{P}^\sigma$ respectively, which are given by:

$$\boldsymbol{P}^\mu = \begin{bmatrix} p^\mu_{ll} & 1-p^\mu_{hh} \\ 1-p^\mu_{ll} & p^\mu_{hh} \end{bmatrix}, \quad \boldsymbol{P}^\sigma = \begin{bmatrix} p^\sigma_{ll} & 1-p^\sigma_{hh} \\ 1-p^\sigma_{ll} & p^\sigma_{hh} \end{bmatrix}. \tag{2}$$

The agent cannot observe the state of the economy and must infer it from realized consumption growth. In contrast to the long-run risk model in Bansal and Yaron (2004), Bayesian learning embedded in this consumption process is able to generate time-series estimates of expected consumption growth and consumption volatility, which are further analyzed in the following empirical asset pricing tests. The posterior belief over specific states at date $t+1$ conditional on observations available until date t is denoted by the vector $\xi_{t+1|t}$. Bayesian inference implies that the belief vector evolves according to:

$$\xi_{t+1|t} = \boldsymbol{P}\frac{\left(\xi_{t|t-1} \odot \boldsymbol{\eta}_t\right)}{\mathbf{1}'\left(\xi_{t|t-1} \odot \boldsymbol{\eta}_t\right)}, \tag{3}$$

where $\boldsymbol{\eta}_t$ is a vector of conditional Gaussian densities, $\odot$ represents element-by-element multiplication, $\boldsymbol{P} = \boldsymbol{P}^\mu \otimes \boldsymbol{P}^\sigma$ is the joint transition matrix, and $\otimes$ is the Kronecker product. Despite the mean and volatility states switching independently, Bayesian learning implies that the agent's beliefs over those states are dependent (Lettau et al. (2008) and Boguth and Kuehn (2013)).

[1] Aretz et al. (2022) use a SDF model to illustrate that expected European option returns are not unambiguously related to their underlying asset's volatility, with the sign of the relation depending on the option's moneyness.

2.2. The stochastic discount factor

The agent's preferences obey Epstein and Zin's (1989) recursive utility function, given by:

$$U_t = \left[(1-\beta) C_t^{1-\frac{1}{\psi}} + \beta \left[E_t \left(U_{t+1}^{1-\gamma} \right) \right]^{\frac{1-\frac{1}{\psi}}{1-\gamma}} \right]^{\frac{1}{1-\frac{1}{\psi}}}, \tag{4}$$

where β is the time discount factor, γ the RRA parameter, ψ the EIS, U_{t+1} the continuation value at time $t+1$, and $\gamma > 0$, $\psi > 0$, and $\psi \neq 1$. For $\psi = \frac{1}{\gamma}$, the representative agent has standard constant relative risk aversion (CRRA) preferences. When $\psi > (<) \frac{1}{\gamma}$, the agent prefers early (late) resolution of uncertainty.

The Euler equation is given by:

$$E_t \left[M_{t+1} R_{i,t+1} \right] = 1, \tag{5}$$

where M_{t+1} is the SDF, and $R_{i,t+1}$ is the return on any asset. The SDF under recursive utility is

$$M_{t+1} = \beta^{\theta} \left(\frac{C_{t+1}}{C_t} \right)^{-\frac{\theta}{\psi}} \left(R_{t+1}^W \right)^{\theta - 1},$$

where $\theta = \frac{1-\gamma}{1-1/\psi}$, and R_{t+1}^W is the return on wealth, given by $R_{t+1}^W = W_{t+1}/(W_t - C_t)$, with W_{t+1} representing wealth at time $t+1$. According to the Euler equation, the log wealth-consumption ratio, which is defined as $z_t = \log(W_t/C_t)$ and depends on the agent's belief $z_t = z\left(\xi_{t+1|t}\right)$, satisfies the functional equation

$$E_t \left[\exp \left(\theta \left(\log \beta + \left(1 - \frac{1}{\psi} \right) \Delta c_{t+1} + z_{t+1} - \log \left(e^{z_t} - 1 \right) \right) \right) \right] = 1.$$

2.3. Asset pricing implications

Although the wealth-consumption ratio is not observable in the data, we can use numerical simulations to show that the change in z_t is approximately linear in the change in the conditional mean of consumption growth, denoted by $\Delta\hat{\mu}_t$, and the change in the conditional volatility of consumption growth, denoted by $\Delta\hat{\sigma}_t$. As a result, we can further linearize the SDF and derive the following pricing equation for any asset i from the Euler equation:

$$E_t \left[R_{i,t+1}^e \right] \approx \beta_{\Delta c,t}^i \lambda_{\Delta c,t} + \beta_{\Delta \mu,t}^i \lambda_{\Delta \mu,t} + \beta_{\Delta \sigma,t}^i \lambda_{\Delta \sigma,t}, \tag{6}$$

where $E_t \left[R_{i,t+1}^e \right]$ is asset i's expected excess return, $\beta_{\Delta c,t}^i$, $\beta_{\Delta \mu,t}^i$, and $\beta_{\Delta \sigma,t}^i$ are the asset's exposures to consumption growth and its conditional mean and volatility respectively, and $\lambda_{\Delta c,t}$, $\lambda_{\Delta \mu,t}$, and $\lambda_{\Delta \sigma,t}$ are the corresponding risk premiums of the three consumption exposures. It is worth noting that the exact functional forms of $\lambda_{\Delta \mu,t}$ and $\lambda_{\Delta \sigma,t}$ are unknown due to the approximation of the change in z_t.

Due to risk aversion, the model always predicts a positive risk premium for the consumption growth exposure. Moreover, it can be shown that when $\psi > 1/\gamma$ (the agent prefers early resolution of uncertainty), the risk premiums on the conditional mean and volatility of consumption growth are positive and negative respectively. Interested readers can refer to Boguth and Kuehn (2013) and the Internet Appendix for the derivation of the above-mentioned results.

3. Impulse response analysis

In this section, we use an impulse response analysis to study how consumption risks affect delta-hedged option returns within the framework presented in Section 2. We then discuss the mechanism of our results.

3.1. Computing model-implied option returns

To compute the delta-hedged return, we begin by considering a stock paying dividends that are positively correlated with aggregate consumption. Specifically, we follow Abel (1999) and Bansal and Yaron (2004) and assume that the dividend growth process, Δd_t, is given by:

$$\Delta d_t \equiv \ln \left(\frac{D_t}{D_{t-1}} \right) = \Phi \Delta c_t + g_d + \sigma_d \epsilon_{d,t} \tag{7}$$

where D_t is the dividend at time t, Φ is the leverage parameter, $\epsilon_{d,t}$ is a standard i.i.d. normal shock independent of other shocks in the model, σ_d is the volatility of that idiosyncratic shock, and g_d is a constant. In equilibrium, the price–dividend ratio $\frac{S_t}{D_t} \equiv \varphi\left(\xi_{t+1|t}\right)$ satisfies the Euler equation:

$$\frac{S_t}{D_t} = E_t \left[M_{t+1} \left(\frac{S_{t+1}}{D_{t+1}} + 1 \right) \frac{D_{t+1}}{D_t} \right]$$

or equivalently,

$$\varphi\left(\xi_{t+1|t}\right)=E_t\left[M_{t+1}\left(\varphi\left(\xi_{t+2|t+1}\right)+1\right)\exp\left(\Delta d_{t+1}\right)\right]. \tag{8}$$

We solve the fixed point of the price-dividend ratio as determined by Eq. (8) using the linear interpolation method.[2] The model-implied risk-free rate is $r_t^f \equiv \ln\left(R_t^f\right)$, with $R_t^f = 1/E_t\left[M_{t+1}\right]$.

We assume that the option price is equal to the option value implied by the equilibrium model. For instance, the current value of a call option expiring in n periods, $\tilde{C}_t^{(n)}$, is given by the expectation of the option's future cash flow multiplied by the multi-period SDF:

$$\tilde{C}_t^{(n)}=E_t\left[M_{t,t+n}\max\left(0,S_{t+n}-K\right)\right], \tag{9}$$

where S_{t+n} is the price of the underlying asset at time $t+n$, K is the option's strike price, and $M_{t,t+n}$ is the multi-period SDF, $M_{t,t+n}=M_{t,t+1}M_{t+1,t+2}\cdots M_{t+n-1,t+n}$ in which $M_{t,t+1}$ is the one-period SDF.

We set the leverage parameter $\Phi=3$, in line with previous studies such as Bansal and Yaron (2004) and Lettau et al. (2008). The parameters g_d and σ_d are set to match the unconditional mean and standard deviation of dividend growth in the post-war data, which yields quarterly values $g_d=-0.009$ and $\sigma_d=0.028$. Our calibration analysis suggests that the benchmark case with ($\gamma=60$, $\psi=1.5$) and our empirical estimates of the Markov-switching model parameters in Table 1 can generate a sizable equity premium of about 4% per year when the model is calibrated at the quarterly frequency. Since ($\gamma=60$, $\psi=1.5$) implies that the agent prefers early resolution of uncertainty, we also run simulations with ($\gamma=60$, $\psi=0.01$) to study the case in which the agent prefers late resolution of uncertainty.

We use Monte Carlo simulations to compute the values of an ATM European call option with three months (twelve weeks) to maturity according to Eq. (9).[3] The computation involves simulating 40,000 sample paths of stock prices and the multi-period SDF using the parameters that match the aggregate consumption process, as shown in Table 1. Along a sample path, we track the contract and compute the option prices given time-t state belief $\xi_{t+1|t}$, the current dividend D_t, and the pre-specified strike price K. We compute the implied volatility and delta derived from the Black–Scholes model for each option that is still alive at time t. Because the underlying asset is a dividend-paying stock, we make appropriate adjustments to the equilibrium price S_t and use the ex-dividend price in computing the implied volatility and delta of the options.

We calculate the model-implied delta-hedged gain of a call option over its lifetime as:

$$\begin{aligned}\Pi(t,t-11) = {} & \tilde{C}_t^{(1)}-\tilde{C}_{t-11}^{(12)}-\sum_{n=0}^{10}\Delta_{c,t-11+n}\left(S_{t-10+n}-S_{t-11+n}\right)\\ & -\sum_{n=0}^{10} r_{t-11+n}^f\left(\tilde{C}_{t-11+n}^{(12-n)}-\Delta_{c,t-11+n}S_{t-11+n}\right),\end{aligned}$$

where $\tilde{C}_{t-11}^{(12)}$ is the value of the option when issued, $\tilde{C}_t^{(1)}$ the option value one period (week) before expiration, and $\Delta_{\tilde{c},t-11+n}$, $\tilde{C}_{t-11+n}^{(12-n)}$, S_{t-11+n}, and r_{t-11+n}^f ($n=0,1,\ldots,10$) are, respectively, the option delta, option value, stock price, and the risk-free rate within the horizon of the delta-hedged gain. The value of the delta-hedged option at the start of the horizon is $\Delta_{\tilde{c},t-11}S_{t-11}-\tilde{C}_{t-11}^{(12)}$. The delta-hedged gain divided by the absolute value of the delta-hedged option portfolio yields the delta-hedged return.

The model-implied delta-hedged return resembles its empirical counterpart. First, the delta-hedged return is an excess return derived from a self-financing strategy. Second, the computation of the delta-hedged gain requires multiple intermediate delta-hedging opportunities within the horizon. Third, the model-implied delta-hedged return crucially depends on stock and option prices determined in equilibrium. As such, we can investigate the mechanism of the model by examining the impact of consumption risks on the SDF and equilibrium asset prices.

3.2. Simulation results

We perform impulse response analysis to study the impacts of changing beliefs about the consumption growth regimes on the SDF, the stock return, the call option return, and the delta-hedged call option return. First, we assume that consumption growth stays at its long-run mean implied by the estimated Markov-switching model. Due to Bayesian learning, the agent's belief converges to the stationary level. We then suppose that a negative shock to consumption growth occurs in the fifth period. The agent updates his belief according to Bayes' rule, leading to a decline in the posterior probability of the high mean growth regime and an increase in the posterior probability of the high volatility regime. Consequently, the agent's estimate of mean growth falls, whereas his estimate of consumption volatility rises. The top two panels of Fig. 1 display these results. Assuming ($\gamma=60$, $\psi=1.5$), the other panels in Fig. 1 present the impulse responses of the SDF, the stock return, the conditional variance of the stock return, the implied volatility, the call option return, and the delta-hedged call option return in response to the negative shock. Under recursive utility,

[2] We choose 50 grid points on each dimension of the state variables and use function iterations to find the fixed points for both the wealth-consumption ratio and the price-dividend ratio.

[3] Before running Monte Carlo simulations to compute delta-hedged returns, we solve the model (including the SDF and price-dividend ratio) numerically at the weekly frequency by appropriately scaling relevant parameters in the model. An alternative approach is to develop a continuous-time asset pricing model with recursive utility and a hidden Markov model. However, (semi)closed-form solutions are not available for such a model. Moreover, because volatility is instantaneously observable in the continuous-time setting, it would be infeasible to analyze the impact of learning about the volatility state on equilibrium prices.

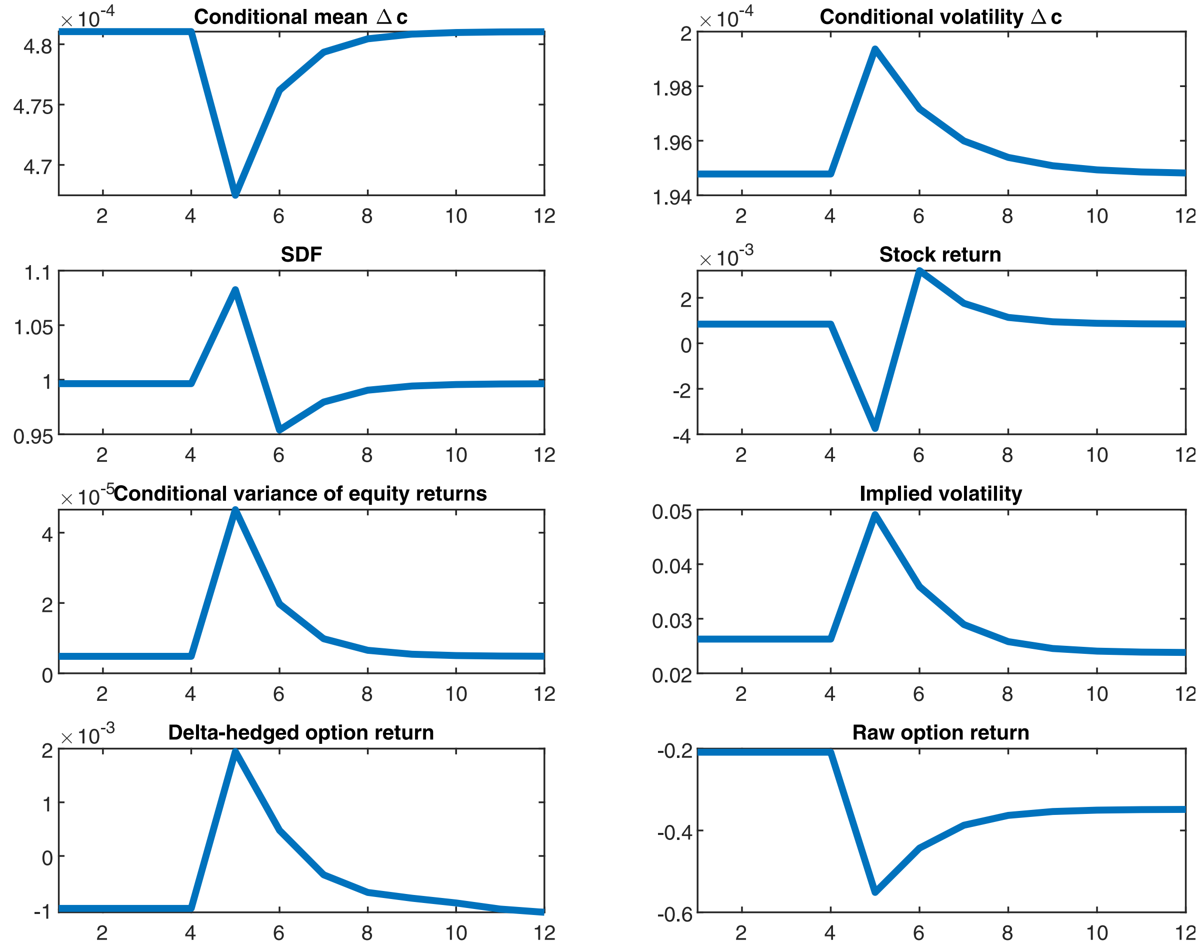

Fig. 1. Impulse responses: Conditional mean and volatility, $\gamma = 60, \psi = 1.5$. This figure plots the impulse response functions when the growth rate of consumption shifts from its long-run mean to the low mean growth rate. The agent's belief vector $\xi_{t+1|t}$ is updated according to Bayes' rule. The risk aversion parameter is set at $\gamma = 60$, and the EIS parameter at $\psi = 1.5$.

the continuation value falls as a result of the lower conditional mean and the higher conditional volatility of consumption growth. Because the values of γ and ψ imply a preference for early resolution of uncertainty, the SDF rises significantly in response to the shock. Moreover, the stock return drops as the equity value depreciates, while the conditional variance of the stock return rises due to an enhanced pessimism about the state of the economy. The co-movement of the SDF and the conditional stock variance suggests that stock return variance carries a risk premium. This is also evident from the observation that the implied volatility of the call option increases substantially. The lowest panel in the figure shows that the call option return falls because the effect of lower equity value dominates that of higher implied volatility. On the contrary, the delta-hedged call option gain (return) rises due to the elimination of the impact of the underlying stock price movement on the call option price.

The co-movement of the SDF and the delta-hedged call option return in Fig. 1 implies that the delta-hedged option enables investors to hedge against systematic risk. Because consumption growth and its conditional mean are negatively related to the SDF, both factors have positive risk premiums. In contrast, since conditional consumption volatility is positively related to the SDF, it has a negative risk premium. These results are consistent with our empirical findings presented in Sections 4.3 and 4.4 below.

Boguth and Kuehn (2013) find that consumption volatility is important to price stocks. To focus on volatility risk on its own, we run a second impulse response analysis by assuming that consumption growth and its expectation remain unchanged but consumption volatility rises. Fig. 2 shows that the responses of the variables of interest are largely similar to those in the previous case. The rise in conditional volatility alone leads to an increase in the SDF, the conditional stock variance and the implied volatility. In addition, the option return falls, whereas the delta-hedged return rises. Thus, the delta-hedged option represents a hedging opportunity against consumption volatility, and consumption volatility carries a negative risk premium.

Figs. 3 and 4 plot simulation results for the ($\gamma = 60$, $\psi = 0.01$) case, in which the agent prefers late resolution of uncertainty. Fig. 3 shows that the delta-hedged return and the SDF move in opposite directions in response to a negative mean growth shock despite that both conditional stock variance and implied volatility rise on impact. Contrary to the ($\gamma = 60$, $\psi = 1.5$) case and our empirical evidence, in this case the delta-hedged call option has a positive exposure to systematic risk in that the return on the

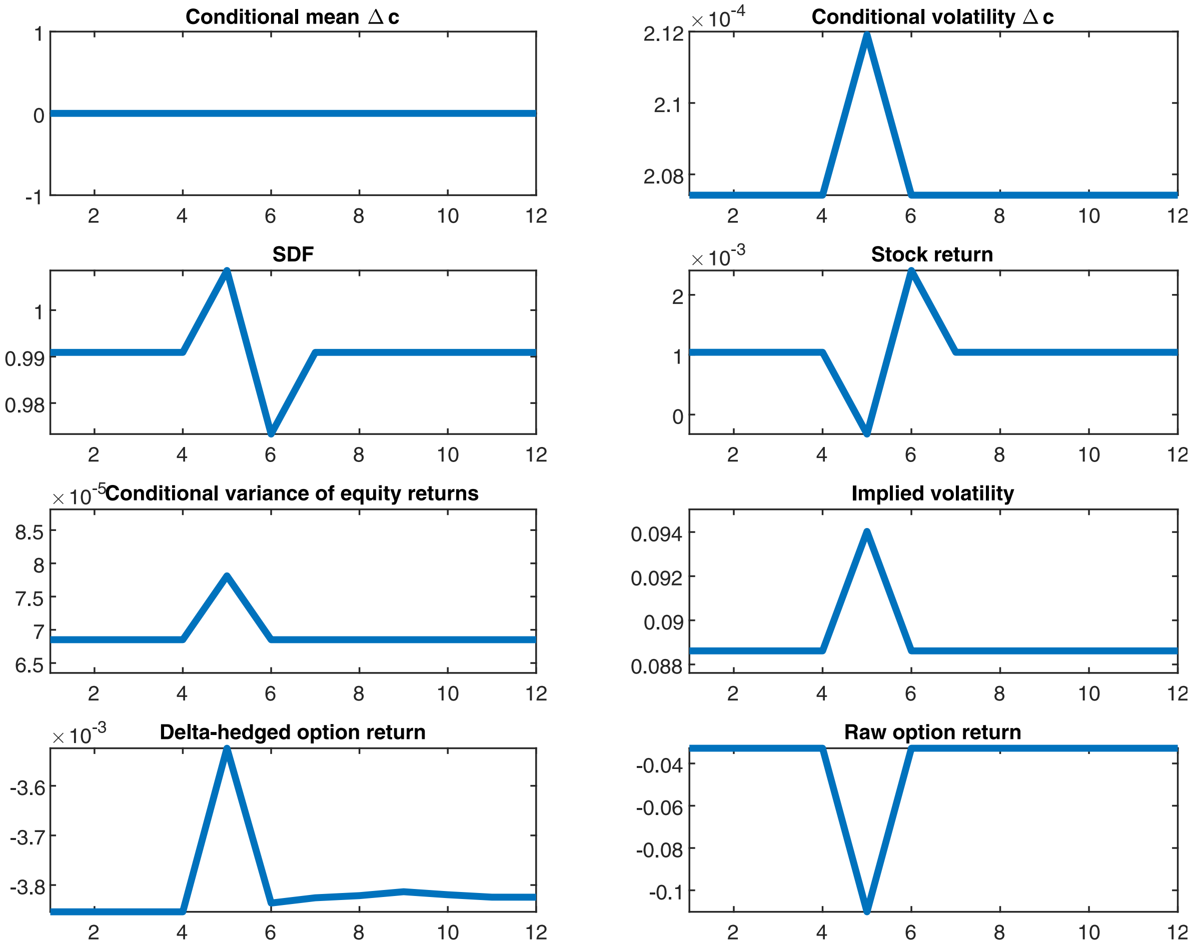

Fig. 2. Impulse responses: Conditional volatility, $\gamma = 60, \psi = 1.5$. This figure plots the impulse response functions when conditional volatility of consumption growth rises. The risk aversion parameter is set at $\gamma = 60$, and the EIS parameter at $\psi = 1.5$.

delta-hedged option performs poorly when the SDF is high. On the other hand, the stock return increases and co-moves with the SDF. Thus, the implied risk premiums on the conditional mean and volatility of consumption, respectively, have opposite signs compared to the ($\gamma = 60$, $\psi = 1.5$) case. As shown in Fig. 4, when the conditional volatility of consumption growth rises on its own, implied volatility falls, resulting in a decline in the delta-hedged option return. Because of the agent's preference for late resolution of uncertainty and the absence of shocks to the level of consumption, the SDF drops in response to the volatility shock. Although the delta-hedged option return and the SDF co-move in the same direction, the implied risk premium on consumption volatility is positive, opposite to that in the ($\gamma = 60$, $\psi = 1.5$) case and in our empirical evidence. In Sections 4.3 and 4.4, our GMM estimation identifies a negative relation between the SDF and the conditional mean of consumption growth while a positive relation between the SDF and consumption volatility.

4. Empirical tests

In this section, we use options data to estimate the risk premiums of consumption growth, mean growth, and consumption volatility. We first fit a Markov-switching model to obtain time-series estimates of the conditional mean and volatility of consumption growth. We next sort single-name options into portfolios and use the portfolios to study the pricing of consumption exposures in option returns. We then study whether consumption risk premiums estimated from option returns can price stocks, and how the three consumption risk exposures price the joint cross-section of stock and option returns.

4.1. Estimating consumption dynamics

Defining total consumption as the sum of non-durable goods consumption expenditures and service consumption expenditures, we obtain quarterly per capita real expenditures data on the two consumption components from the Bureau of Economic Analysis (BEA). Our sample is from 1952:Q1 to 2018:Q1.

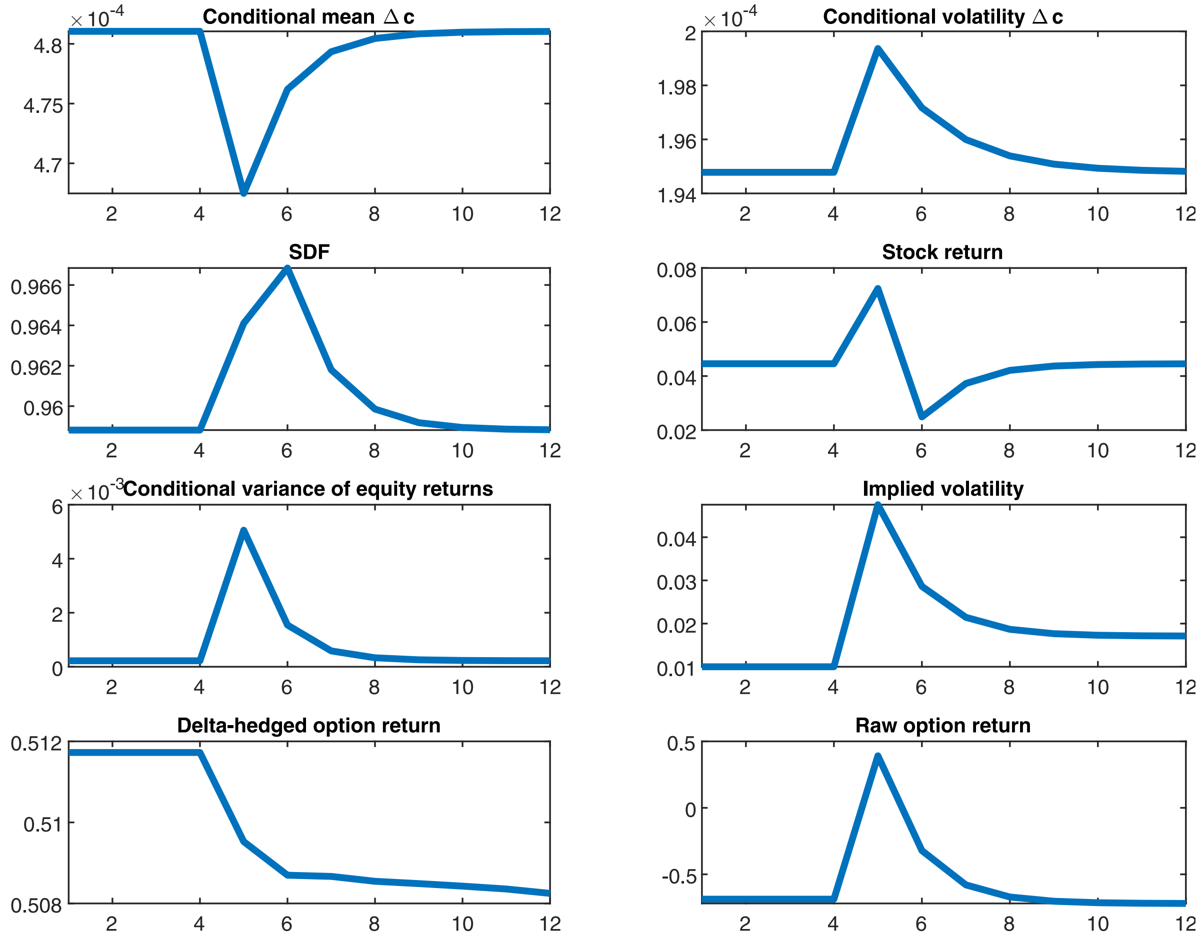

Fig. 3. Impulse responses: Conditional mean and volatility, $\gamma = 60, \psi = 0.01$. This figure plots the impulse response functions when the growth rate of consumption shifts from its long-run mean to the low mean growth rate. The agent's belief vector $\xi_{t+1|t}$ is updated according to Bayes' rule. The risk aversion parameter is set at $\gamma = 60$, and the EIS parameter at $\psi = 0.01$.

We fit a four-state Markov-switching model to our consumption data. While maintaining the assumption that the agent has preferences over total consumption, we follow Boguth and Kuehn (2013) in separately using non-durable and service consumption expenditures in our estimation to improve state identification and to reduce standard errors. In particular, we assume that both log non-durable goods consumption growth and the log change in the share of non-durable to total consumption follow Markov chains. This strategy implies that log total consumption growth also follows a Markov chain. More precisely, we express total consumption C_t as non-durable goods consumption N_t divided by the non-durable consumption share V_t. That is, $C_t = N_t/V_t$. Thus, log total consumption growth is log non-durable consumption growth, Δn_t, minus the log change in the non-durable consumption share, Δv_t:

$$\Delta c_{t+1} = \Delta n_{t+1} - \Delta v_{t+1}. \tag{10}$$

Given Eq. (10), we assume that both Δn_{t+1} and Δv_{t+1} follow Markov chains:

$$\Delta n_{t+1} = \mu_t^n + \sigma_t^n \epsilon_{t+1}^n \quad \Delta v_{t+1} = \mu_t^v + \sigma_t^v \epsilon_{t+1}^v, \tag{11}$$

where μ_t^k and σ_t^k, with $k \in \{n, v\}$, are, respectively, the mean of log non-durable consumption growth ($k = n$) or the log change in the non-durable consumption share ($k = v$) and the volatility of the individual innovation. The innovation ϵ_{t+1}^k is standard normal, with $\mathrm{Cov}_t\left(\epsilon_{t+1}^n, \epsilon_{t+1}^v\right) = \rho_{nv}$. The dynamics specified in Eqs. (10) and (11) together with the fact that the information set $\mathcal{F}_t$ contains $\{\Delta n_t, \Delta v_t\}$ and its history also imply a Markov process for total consumption growth, with dynamics specified in Eq. (1) and $\mu_t = \mu_t^n - \mu_t^v$ and $\sigma_t^2 = \left(\sigma_t^n\right)^2 + \left(\sigma_t^v\right)^2 - 2\rho_{nv}\sigma_t^n\sigma_t^v$. Thus, the estimates of the parameters $\{\mu_h^n, \mu_l^n, \mu_h^v, \mu_l^v, \sigma_h^n, \sigma_l^n, \sigma_h^v, \sigma_l^v, \rho_{nv}\}$ allow us to recover the dynamics of log total consumption growth, which we use to solve the consumption-based model in Section 2.

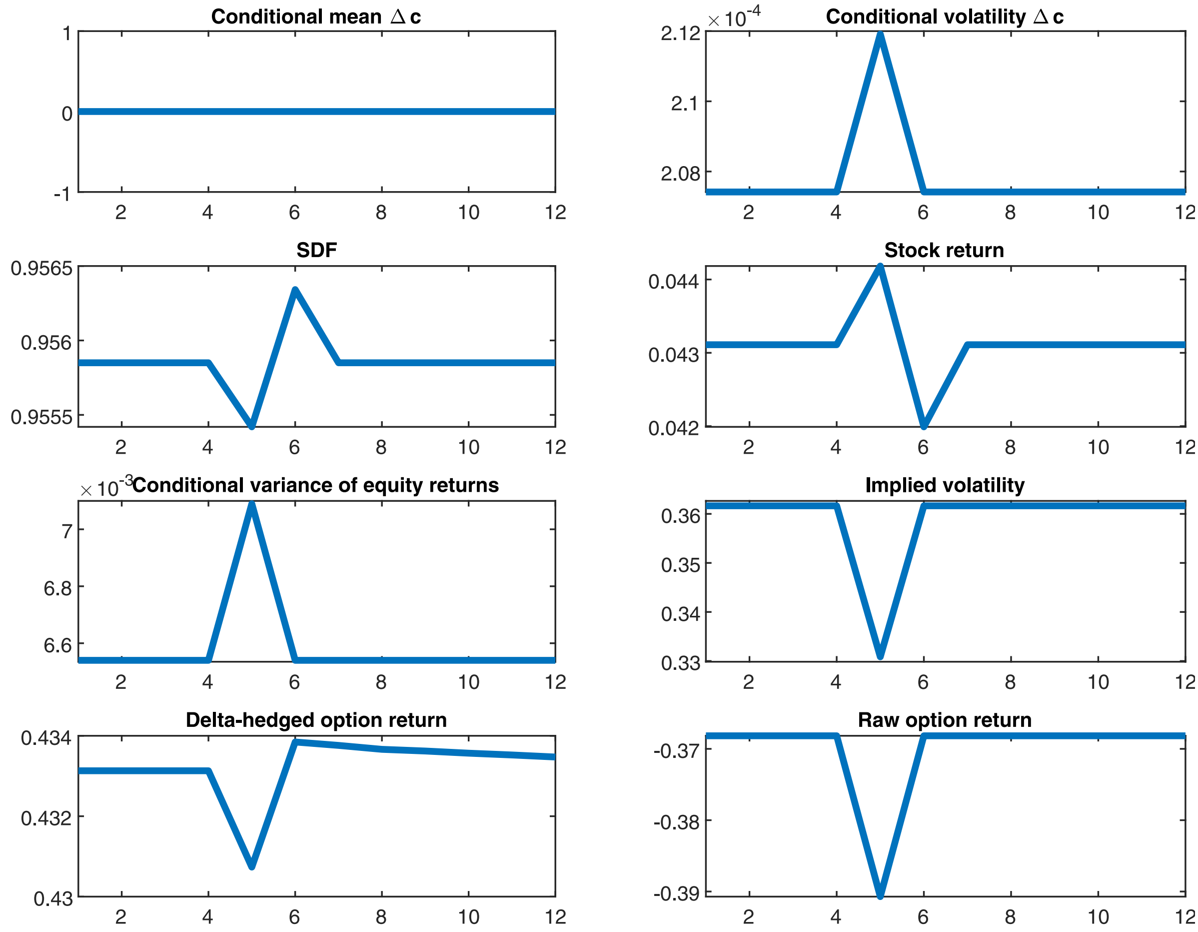

Fig. 4. Impulse responses: Conditional volatility, $\gamma = 60, \psi = 0.01$. This figure plots the impulse response functions when conditional volatility of consumption growth rises. The risk aversion parameter is set at $\gamma = 60$, and the EIS parameter at $\psi = 0.01$.

Table 1 presents the estimates of the parameters in the Markov-switching processes. Panel A shows that expected non-durable consumption growth is positive in the high state ($\mu_h^n = 0.58\%$) and negative in the low state ($\mu_l^n = -0.03\%$). State-conditional non-durable consumption volatilities are $\sigma_l^n = 0.40\%$ and $\sigma_h^n = 0.83\%$. The estimated parameters for the non-durable consumption share (shown in Panel B) are $\mu_l^v = -0.16\%$ and $\mu_h^v = 0.00\%$ and $\sigma_l^v = 0.34\%$ and $\sigma_h^v = 0.58\%$. The correlation between log changes in the two variables is $\rho_{nv} = 0.83$. Turning to total consumption growth, its expected growth is $\mu_l = 0.13\%$ in the low state and $\mu_h = 0.58\%$ in the high state, while its volatility is $\sigma_l = 0.23\%$ in the low state and $\sigma_h = 0.50\%$ in the high state.

The two transition probabilities for the mean growth regimes, p_μ^{ll} and p_μ^{hh}, are 0.87 and 0.95 respectively. Consistent with Lettau et al. (2008), the high mean state is thus markedly more persistent than the low mean state. Both volatility states are persistent, with transition probabilities being about 0.91. Given the differences in the sample and in the consumption measure across the two papers, these estimates differ moderately from those in Lettau et al. (2008), who find the volatility states to be even more persistent.

Fig. 5 presents the filtered beliefs for the regimes. The upper panel depicts the belief dynamics for the high mean growth regime, and the lower panel for the high volatility regime. The gray bars in the graphs indicate economic recession periods defined by the National Bureau of Economic Research. The figure further suggests that the low volatility regime becomes more prevalent from 1990 on, as also observed by Kim and Nelson (1999) and Boguth and Kuehn (2013). Despite that, consumption volatility appears to have returned to the high regime during both the 2000–2001 dot-com crash and the 2008–2009 global financial crisis in the post-1990 period. When the economy is in a recession, the probability of being in the high mean state is low, while the probability of being in the high volatility state tends to be high in certain periods. However, their correlation is far from being perfect due to the assumption of independent switching between mean regimes and volatility regimes. Overall, the Markov-switching model captures most recessions during our sample period, with the exceptions of the mild 1969–1970 recession and the 1981–1982 recession caused by contractionary monetary policy.

Table 1
Markov-switching model of consumption growth.

Panel A: Non-durable consumption (%)			
μ_l^n	μ_h^n	σ_l^n	σ_h^n
−0.0269	0.5835	0.4002	0.8255
(−0.57)	(21.51)	(17.67)	(45.73)
Panel B: Non-durable consumption share (%)			
μ_l^v	μ_h^v	σ_l^v	σ_h^v
−0.1576	0.0000	0.3422	0.5835
(−4.46)	(0.00)	(17.63)	(22.01)
Panel C: Marginal transition probabilities			
p_μ^{ll}	p_μ^{hh}	p_σ^{ll}	p_σ^{hh}
0.87	0.95	0.91	0.91
(18.96)	(47.59)	(28.87)	(24.46)
Panel D: Correlation			
ρ_{nv}			
0.8256			
(45.73)			

This table reports parameter estimates for the Markov-switching model fitting log non-durable goods consumption growth, Δn_{t+1}, and changes in the log non-durable consumption share, Δv_t,

$$\Delta n_{t+1} = \mu_t^n + \sigma_t^n \epsilon_{t+1}^n \quad \Delta v_{t+1} = \mu_t^v + \sigma_t^v \epsilon_{t+1}^v,$$

where for $i \in \{n, v\}$, μ_t^i denotes the mean regime, σ_t^i denotes the volatility regime, and ϵ_{t+1}^i is standard normal with $\mathrm{Cov}_t\left(\epsilon_{t+1}^n, \epsilon_{t+1}^v\right) = \rho_{nv}$. The transition matrices, $\boldsymbol{P}^\mu$ and $\boldsymbol{P}^\sigma$, are given by

$$\boldsymbol{P}^\mu = \begin{bmatrix} p_\mu^{ll} & 1-p_\mu^{hh} \\ 1-p_\mu^{ll} & p_\mu^{hh} \end{bmatrix} \quad \boldsymbol{P}^\sigma = \begin{bmatrix} p_\sigma^{ll} & 1-p_\sigma^{hh} \\ 1-p_\sigma^{ll} & p_\sigma^{hh} \end{bmatrix}.$$

The estimation procedure follows Hamilton (1994). The data for estimation include quarterly per capita real consumption expenditures for non-durable goods and services for the period 1952.Q1 to 2018.Q1. t-statistics are reported in parentheses.

4.2. Calculation of delta-hedged call option returns

We obtain call options data over the period from January 1996 to December 2017 from Optionmetrics.[4] The data include the daily closing bid and ask quotes, the trading volume, the strike price, and the maturity date of each option. The data further include each option's delta, calculated by Optionmetrics using standard market conventions, the closing price of and the dividends paid out by the stocks underlying the options, and the risk-free rate of return.

We apply standard filters to the options data (see Goyal and Saretto (2009) and Cao and Han (2013)). First, we exclude an option if the stock underlying the option pays out dividends over the option's remaining time-to-maturity. Second, we exclude option observations violating well-known arbitrage bounds. More specifically, we exclude an option observation if the option's price does not fulfill $S \geq \tilde{C} \geq \max(0, S - Ke^{-rT})$ where $\tilde{C}$ is the call option's price, S the underlying stock's price, K the strike price, T the option's time-to-maturity, and r the risk-free rate of return in the data. Third, we only retain option observations with positive trading volume, a positive bid quote, a bid price strictly smaller than the ask price, and a bid-ask midpoint of at least \$1/8. Finally, we only keep option observations whose last trade date matches the record date and whose option price date matches the underlying stock's price date.

We use quarterly delta-hedged option returns in our main tests, and monthly delta-hedged option returns in robustness tests. In either case, we calculate the return from the start of a calendar quarter or month to its end. In line with Bakshi and Kapadia (2003) and Cao and Han (2013), we define the delta-hedged option return as the delta-hedged option gain over the period scaled by the absolute value of the delta-hedged option at the start of the period, where the delta-hedged option is a self-financing portfolio consisting of a long option, a hedging position in the underlying stock, and a money market investment. The value of a perfectly delta-hedged option would be insensitive to changes in the value of the underlying stock. Assuming that the delta-hedge is rebalanced at the end of every trading day, we calculate the delta-hedged call option gain over the quarter or month starting at time $t-1$ and ending at time t, $\Pi(t-1, t)$, as:

$$\Pi(t-1, t) = \tilde{C}_t - \tilde{C}_{t-1} - \sum_{n=0}^{N-1} \Delta_{\tilde{c}, t_n} \left[S\left(t_{n+1}\right) - S\left(t_n\right) \right] - \sum_{n=0}^{N-1} \frac{a_n r_{t_n}}{365} \left[\tilde{C}\left(t_n\right) - \Delta_{\tilde{c}, t_n} S\left(t_n\right) \right], \tag{12}$$

where $\tilde{C}_t$ is the call option price at time t, $\Delta_{\tilde{c}, t_n}$ the option delta, r_{t_n} the annualized risk-free rate of return, $S(t_n)$ the underlying stock price at the end of trading day t_n, where $t_n \in \{t_0, t_1, \ldots, t_{N-1}\}$ are the N trading days within the period from time $t-1$ to t, and a_n is the number of calendar days between t_n and t_{n+1}. We finally calculate the value of the delta-hedged call option at time $t-1$ as the absolute value of $\Delta_{\tilde{c}, t-1} S_{t-1} - \tilde{C}_{t-1}$.

[4] We only consider American call options written on non-dividend stocks in our empirical tests since American put options contain an early exercise risk premium correlated with moneyness (see, e.g., Aretz and Gazi (2021)).

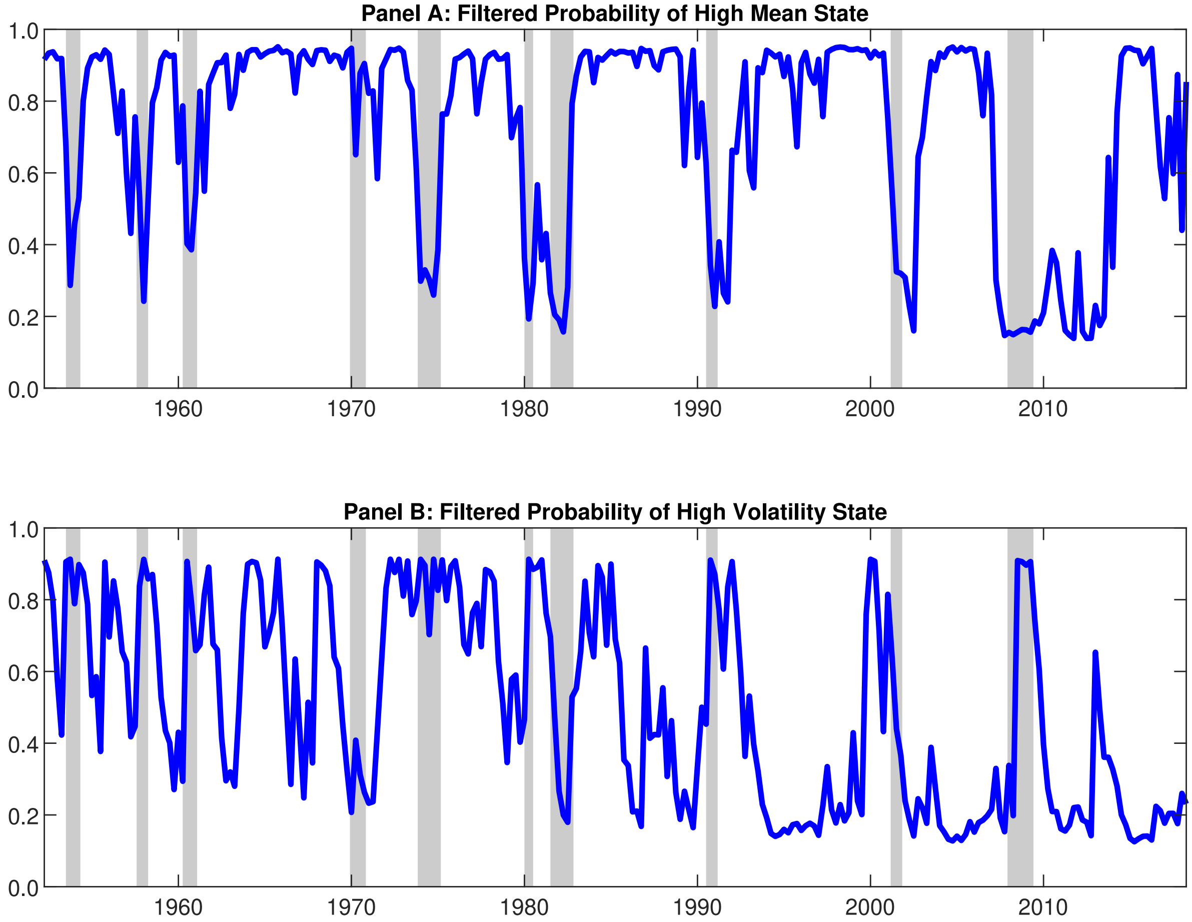

Fig. 5. Bayesian beliefs about the mean and volatility state. This figure displays the estimated Bayesian belief processes for being in the high expected growth rate state (top figure) and high volatility state (bottom figure). The estimation procedure follows Hamilton (1994). We use quarterly per capita real consumption expenditure for non-durable goods and services for the period 1952:Q1–2018:Q1. The parameter estimates for the Markov-switching model are reported in Table 1. The gray bars indicate NBER recession periods.

We select options with a time-to-maturity between 106 to 176 days so that the option expires after the end of the following quarter and the time-to-maturity range is not too wide. Within the maturity range, we select the option with the shortest time-to-maturity, the largest open interest, and the highest trading volume for each stock. The selection results in 28,064 quarterly observations for call options. Table 2 shows summary statistics for the quarterly option sample. The average delta-hedged call option return is −2.47% per quarter, with a variation of 10.19%. The average call option has a time-to-maturity of 137 days, a daily IVOL of 2.95%, and moneyness of 0.97.

4.3. The pricing of consumption risks in option returns

In this section, we study how consumption growth, expected consumption growth, and consumption volatility exposures price the cross-section of call option returns. We start our investigation by looking at the cross-section of quarterly returns on IVOL and moneyness sorted portfolios of delta-hedged calls.[5]

Cao and Han (2013) discover a strongly significant negative relation between delta-hedged stock-option returns and IVOL. Bakshi and Kapadia (2003) show that delta-hedged call option returns are negative and monotonically increase with options moneyness.[6] Boyer and Vorkink (2014) offer evidence that a call option's ex-ante skewness is negatively related to its return and monotonically decreases with options moneyness. Taken together, these two findings imply that lower-moneyness (OTM) call options

[5] We also consider the realized and implied variance of these portfolios over the same return horizon. As robustness tests, we next analyze the monthly returns of the IVOL and moneyness sorted call option portfolios. The relevant results are presented in the Internet Appendix.

[6] Bakshi and Kapadia (2003) also present evidence that the dollar gains (losses) are smaller for OTM than ITM options. This is to be expected because the dollar gains are not normalized by price and OTM options are much cheaper than ITM options. Once the gains are normalized, OTM options have more negative returns than ITM options.

Table 2
Summary statistics for the call options sample.

	Mean	Median	SD	5th	10th	25th	75th	90th	95th
Delta-hedged returns	−2.47	−2.03	10.19	−19.34	−13.36	−6.61	2.10	7.88	12.89
Days-to-maturity	137	140	26	108	108	112	169	172	173
Idiosyncratic volatility	2.95	2.55	1.90	1.05	1.29	1.78	3.64	5.02	6.16
Moneyness	0.97	0.95	0.19	0.72	0.77	0.86	1.03	1.16	1.27
Bid–ask spread	0.28	0.20	0.33	0.05	0.06	0.10	0.30	0.50	0.80
Volume	96	14	530	1	2	5	49	150	313
Open interest	2719	612	8430	27	59	187	2019	5898	11 347

This table reports descriptive statistics on the delta-hedged call option returns, days-to-maturity, idiosyncratic underlying-stock volatility, moneyness, bid-ask spread, trading volume, and open interest of option contracts sampled at a quarterly frequency. We exclude the following option observations: the stock underlying the option pays out cash over the option's remaining time-to-maturity; option price violates arbitrage bounds; reported trading volume is 0; option bid quote is 0 or midpoint of bid and ask quotes is less than \$1/8. For each optionable stock, we select that call having the shortest time-to-maturity, the largest open interest and the largest trading volume. Then we only keep calls with days-to-maturity within the range from 106 to 176 days. Delta-hedged returns (in percentage) are calculated through option delta-hedged gains (given by Eq. (12)) scaled by the absolute value of $\Delta S - \tilde{C}$, where Δ is the Black–Scholes option delta, S is the underlying stock price, and $\tilde{C}$ is the price of the call option at the beginning of a quarter. Days-to-maturity is the number of calendar days until option expiration. Idiosyncratic volatility (in percentage) is the standard deviation of the residuals of the Fama–French 3-factor model estimated using the daily stock returns over the previous quarter. Moneyness is the ratio of stock price over option strike price. Bid-ask spread is the spread between the lowest closing ask and the highest closing bid. "SD" represents standard deviation. "5th" to "95th" are the 5th to 95th percentiles respectively. The option sample period is from January 1996 to December 2017, and the number of sample observations is 28,064.

have higher ex-ante skewness and earn more negative returns than higher-moneyness (ITM) call options do. Our aim in this section is to find out whether consumption risks can help us understand the negative relation between delta-hedged call option returns and IVOL and why delta-hedged call option returns are more negative for OTM than for ITM options.

At the end of each quarter $t-1$ in our sample period, we sort the delta-hedged call options into 16 portfolios independently according to options moneyness and the IVOL of the underlying stock. We construct equally-weighted portfolios and hold them over quarter t. IVOL is the standard deviation of the residual with respect to the Fama–French three-factor model estimated using daily stock returns over the previous quarter. Moneyness is defined as the ratio of the underlying stock's price to the option's strike price. Panel A of Table 3 reports the average option returns for each portfolio. We form four quartile groups based on the IVOL of the underlying stock and four quartile groups based on options moneyness. The bottom IVOL and bottom moneyness portfolio contains options with IVOL and moneyness in the lowest quartile, while the top IVOL and top moneyness portfolio contains options with IVOL and moneyness in the highest quartile. We also create spread portfolios long portfolios with the highest IVOL and short portfolios with the lowest IVOL while keeping the range of options moneyness unchanged ("HIVOL–LIVOL"). Similarly, we construct spread portfolios long portfolios with the highest options moneyness and short portfolios with the lowest while keeping the IVOL of the underlying stock unchanged.

Results in Table 3 support Cao and Han (2013) in showing that the delta-hedged option returns are all negative and become monotonically more negative with increasing IVOL. In accordance, average returns of the "HIVOL–LIVOL" portfolios are all significantly negative. Specifically, the mean return of the "HIVOL–LIVOL" portfolio with moneyness in the lowest quartile is −3.49% (t = −5.60) per quarter, and the mean return of the "HIVOL–LIVOL" portfolio with moneyness in the highest quartile is −1.64% (t = −4.69) per quarter. In line with Bakshi and Kapadia (2003), the mean return becomes more negative the lower the moneyness of the options in a portfolio is (i.e., the more the options are OTM). In particular, the spread return of portfolio "HMON–LMON" with IVOL in the highest quartile is 4.35% (t = 9.04), and the average return of spread portfolio "HMON–LMON" with IVOL in the lowest quartile is 2.50% (t = 3.52).

We also find that these option portfolio returns are close to being normally distributed, alleviating the concern that non-normality could distort our statistical inference.[7] Related results are reported in the Internet Appendix. While the low moneyness options are more volatile than the high moneyness options, skewness monotonically declines with moneyness, with OTM options being positively skewed and ITM options negatively skewed. Thus, sorting options into portfolios according to moneyness is akin to sorting them into portfolios according to Boyer and Vorkink's (2014) ex-ante skewness proxy derived from the Black–Scholes model, even though we do not explicitly use skewness in our sorts. We also differ with Boyer and Vorkink (2014) in that they focus on raw option returns while we examine delta-hedged returns. The return spread between portfolios with high and low skewness further suggests that investors pay more for the tail probability (lottery feature) in OTM options even after controlling for the directional movement of the underlying asset.

Motivated by the model in Section 2 (see Eq. (6)), we run time-series regressions of option portfolio returns on consumption growth, Δc_t, the change in the conditional mean of consumption growth, $\Delta\hat{\mu}_t$, and the change in consumption volatility, $\Delta\hat{\sigma}_t$:

$$R_t^i = \alpha_t^i + \beta_{\Delta c,t}^i \Delta c_t + \beta_{\Delta\mu,t}^i \Delta\hat{\mu}_t + \beta_{\Delta\sigma,t}^i \Delta\hat{\sigma}_t + \epsilon_t^i, \tag{13}$$

where R_t^i is the delta-hedged quarterly return of option portfolio i over period t, α_t^i is a constant, $\beta_{\Delta c,t}^i$, $\beta_{\Delta\mu,t}^i$, and $\beta_{\Delta\sigma,t}^i$ are exposures to consumption risks, and ϵ_t^i is the residual. We obtain the conditional mean and volatility of consumption growth ($\hat{\mu}_t$ and $\hat{\sigma}_t$,

[7] Consistent with the mild skewness and excess kurtosis of the portfolio returns, we have found that bootstrap inference levels used in either our portfolio sorts or FM regressions are similar to asymptotic inference levels. For the sake of brevity, we do not report the bootstrap inference levels in the paper.

Table 3
IVOL and moneyness sorted option portfolios.

	LMON			HMON	HMON–LMON
Panel A: Average return (%)					
LIVOL	−2.78	−1.36	−0.88	−0.27	2.50
	(−3.49)	(−3.08)	(−2.69)	(−0.87)	(3.52)
	−2.88	−1.87	−0.96	−0.52	2.36
	(−4.03)	(−4.45)	(−2.35)	(−1.31)	(3.97)
	−4.01	−2.29	−1.78	−1.03	2.98
	(−6.58)	(−5.01)	(−3.61)	(−2.79)	(6.27)
HIVOL	−6.26	−3.56	−3.06	−1.91	4.35
	(−8.66)	(−6.36)	(−6.20)	(−4.14)	(9.04)
HIVOL–LIVOL	−3.49	−2.20	−2.18	−1.64	
	(−5.60)	(−5.26)	(−6.92)	(−4.69)	
Panel B: $\hat{\beta}_{\Delta c}$					
LIVOL	−2.35	−0.30	−0.25	0.69	
	−1.50	−0.60	−0.01	0.83	
	−4.74	−3.42	−0.06	−0.72	
HIVOL	−4.17	−2.87	−1.84	−1.40	
Panel C: $\hat{\beta}_{\Delta\mu}$					
LIVOL	−5.98	−0.01	−1.00	−2.72	
	7.89	0.35	0.99	−3.32	
	5.32	0.15	−4.30	−1.70	
HIVOL	5.25	4.67	0.09	−8.90	
Panel D: $\hat{\beta}_{\Delta\sigma}$					
LIVOL	21.27	22.48	15.18	10.86	
	22.89	12.02	13.31	18.28	
	13.26	13.14	13.55	12.10	
HIVOL	31.15	26.64	33.74	7.39	

This table reports characteristics of equally-weighted call option portfolios independently sorted on options moneyness (MON) and idiosyncratic underlying-stock volatility (IVOL). Four quartile groups are formed based on the IVOL of the underlying stock and four quartile groups based on options moneyness at the end of quarter $t-1$. The bottom IVOL and bottom moneyness portfolio contains options with IVOL and moneyness in the lowest quartile, while the top IVOL and top moneyness portfolio contains options with IVOL and moneyness in the highest quartile. Portfolios are held over quarter t and rebalanced every quarter. The average return of each portfolio is reported in Panel A. Column "HIVOL–LIVOL" shows the average return of the long-short strategy which buys the highest IVOL portfolio and sells the lowest IVOL portfolio. "HMON–LMON" presents the average return of the portfolio strategy that longs the highest moneyness (in-the-money) portfolio and shorts the lowest moneyness (out-of-the-money) portfolio. Full sample loadings on consumption growth ($\hat{\beta}_{\Delta c}$), the change in conditional mean of consumption growth ($\hat{\beta}_{\Delta\mu}$), and the change in consumption growth volatility ($\hat{\beta}_{\Delta\sigma}$) are reported for each option portfolio in Panels B to D respectively. Newey–West (Newey and West, 1987) adjusted t-statistics with a lag of four are reported in parentheses. The sample period is from January 1996 to December 2017.

respectively) from the estimates of the Markov-switching model in Section 4.1. We estimate regression model (13) over rolling windows spanning ten years of quarterly data, expanding the rolling windows on a quarterly basis. The first rolling window stretches from the second quarter of 1996 to the first quarter of 2006.

Panels B, C, and D of Table 3 present the full-sample exposures of the call option portfolios. We calculate the exposures for IVOL-sorted portfolios by averaging their full sample exposures across moneyness quartiles and those for moneyness-sorted portfolios by averaging the full sample exposures across IVOL quartiles. As displayed in Fig. 6, both IVOL-sorted and moneyness-sorted portfolios produce negative consumption growth exposures, with the exposures becoming more negative over the IVOL-sorted portfolios and less negative over the moneyness-sorted portfolios. Negative consumption growth exposures suggest that delta-hedged call options are countercyclical assets. Conversely, the same portfolios produce positive consumption volatility exposures, with those exposures becoming more positive over the IVOL-sorted portfolios and less positive over the moneyness-sorted portfolios. There is no clear trend in the mean growth exposures over the IVOL-sorted portfolios and a decreasing trend over the moneyness-sorted portfolios. That the high IVOL portfolios and the low moneyness (OTM) portfolios are more negatively exposed to consumption growth and more positively to consumption volatility suggests that they are better suited to hedge against consumption risks than the low IVOL portfolios and the high moneyness (ITM) portfolios. The only concerning aspect of the low-moneyness portfolios is that they can have a higher mean growth exposure than the other portfolios, which lowers their ability to hedge against adverse economic conditions.

To test whether the three consumption exposures are priced, we next run FM regressions of the quarterly returns of the option portfolios on subsets of the exposures. In our most comprehensive specification, we regress the quarterly return of option portfolio i over quarter $t+1$, R^i_{t+1}, on the consumption growth exposure, $\hat{\beta}^i_{\Delta c,t}$, the mean growth exposure, $\hat{\beta}^i_{\Delta\mu,t}$, and the consumption volatility exposure, $\hat{\beta}^i_{\Delta\sigma,t}$, of the portfolio:

$$R^i_{t+1} = \varphi_{0,t+1} + \lambda_{\Delta c,t+1}\hat{\beta}^i_{\Delta c,t} + \lambda_{\Delta\mu,t+1}\hat{\beta}^i_{\Delta\mu,t} + \lambda_{\Delta\sigma,t+1}\hat{\beta}^i_{\Delta\sigma,t} + \eta^i_{t+1}, \tag{14}$$

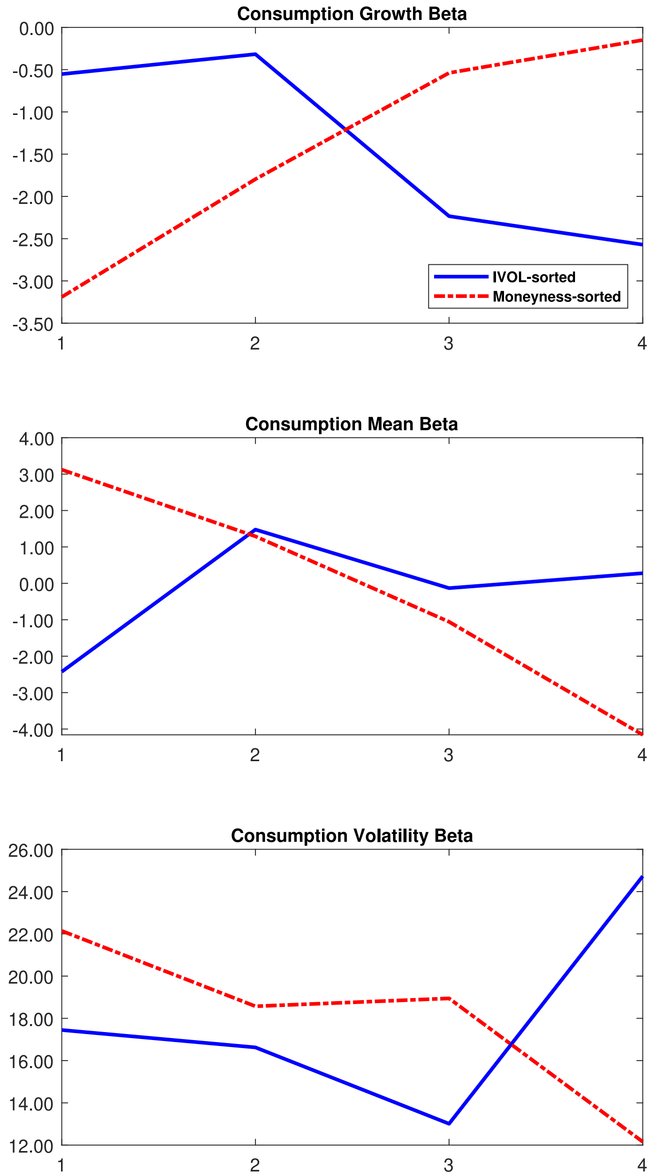

Fig. 6. Consumption betas of IVOL and moneyness sorted option portfolios. This graph displays consumption risk exposures estimated using the full sample for IVOL-sorted and moneyness-sorted portfolios. The exposures for IVOL-sorted portfolios are calculated by averaging the full sample exposures across moneyness quartiles and those for moneyness-sorted portfolios by averaging the full sample exposures across IVOL quartiles. The exposures for IVOL-sorted portfolios are plotted with the blue line, and those for moneyness-sorted portfolios with the red dashed line. Portfolio 1 represents the IVOL-sorted portfolio with idiosyncratic underlying-stock volatility in the lowest quartile, and portfolio 4 in the highest quartile. Meanwhile, portfolio 1 indicates the moneyness-sorted portfolio with options moneyness in the lowest quartile, and portfolio 4 in the highest quartile. The upper panel, middle panel, and lower panel show consumption growth beta ($\hat{\beta}_{\Delta c}$), consumption mean beta ($\hat{\beta}_{\Delta \mu}$), and consumption volatility beta ($\hat{\beta}_{\Delta \sigma}$) respectively. The sample period is from January 1996 to December 2017.

where $\varphi_{0,t+1}$ is a constant and $\lambda_{\Delta c,t+1}$, $\lambda_{\Delta \mu,t+1}$, and $\lambda_{\Delta \sigma,t+1}$ are the risk premiums. The exposure estimates, $\hat{\beta}^i_{\Delta c,t}$, $\hat{\beta}^i_{\Delta \mu,t}$, and $\hat{\beta}^i_{\Delta \sigma,t}$, are obtained from the rolling-window time-series regressions.

Table 4 Panel A presents the FM regression results. Model specification I tests the standard CCAPM. The risk premium estimate, $\lambda_{\Delta c}$, is significant at 0.62% (t = 5.83) and the average cross-sectional R^2 is around 27.64%. Prior studies (e.g., Mankiw and Shapiro (1986), Lettau and Ludvigson (2001), and Boguth and Kuehn (2013)) show that quarterly contemporaneous consumption growth exposures do not explain stock returns. Different from prior studies using stock returns, we find that consumption growth exposures significantly explain delta-hedged option returns. Specification II further incorporates the mean consumption growth exposure. As such, the $\lambda_{\Delta c}$ estimate is hardly affected, the risk premium estimate on the mean growth exposure, $\lambda_{\Delta \mu}$, is −0.01% (t = −0.47), and the average R^2 increases to 35.49%. Specification III has the consumption volatility exposure in addition to the consumption growth exposure. Again, the $\lambda_{\Delta c}$ estimate is hardly affected. However, the $\lambda_{\Delta \sigma}$ estimate is −0.05% and significant (t = −3.71), and the

Table 4
Estimating risk premiums with IVOL and moneyness sorted option portfolios.

Panel A: Fama–MacBeth regressions

	Intercept	$\beta^i_{\Delta c,t}$	$\beta^i_{\Delta\mu,t}$	$\beta^i_{\Delta\sigma,t}$	Avg. R^2
I	−0.34	0.62			27.64
	(−0.73)	(5.83)			
II	−0.13	0.58	−0.01		35.49
	(−0.30)	(5.51)	(−0.47)		
III	0.23	0.59		−0.05	37.43
	(0.73)	(5.24)		(−3.71)	
IV	0.34	0.57	0.01	−0.05	45.27
	(1.10)	(4.69)	(0.34)	(−3.45)	

Panel B: GMM tests

	Δc	$\Delta\mu$	$\Delta\sigma$	MAE	RMSE	J	R^2
b	2.49			0.26	0.31	10.82	97.62
	(9.54)					(0.77)	
λ	0.32						
	(9.54)						
b	2.24	−9.70		0.11	0.15	14.07	99.46
	(9.59)	(−3.92)				(0.44)	
λ	0.20	−0.02					
	(5.09)	(−1.36)					
b	2.76		−26.69	0.21	0.23	11.02	98.65
	(11.15)		(−4.97)			(0.68)	
λ	0.37		−0.05				
	(11.14)		(−5.12)				
b	2.26	−9.37	−1.71	0.11	0.15	13.17	99.47
	(9.26)	(−2.57)	(−2.40)			(0.43)	
λ	0.21	−0.02	−0.004				
	(5.11)	(−0.61)	(−2.08)				

This table reports quarterly consumption risk premium estimates using IVOL-moneyness sorted call option portfolio returns. In Panel A, we report the Fama–MacBeth regression results. The consumption risk exposures are estimated from ten-year rolling window time-series regressions (see Eq. (13)). In the cross-section, we regress quarterly option returns over quarter $t+1$ on estimated beta loadings (see Eq. (14)). Panel B reports the GMM estimates of the moment conditions in Eq. (15), showing both the b estimates as well as the implied risk premiums (λ). MAE and RMSE (in percentage) refer to the mean absolute pricing error and the root mean squared error, respectively. Newey–West (Newey and West, 1987) adjusted t-statistics using four lags are reported in parentheses, and p-values for J-statistics are shown in parentheses below the associated J-statistics. The sample period is from January 1996 to December 2017.

average R^2 increases to 37.43%. The negative $\lambda_{\Delta\sigma}$ estimate, which is similar to the estimate in Boguth and Kuehn (2013), identifies a channel for macroeconomic volatility to be priced in options.

Specification IV presents the full three-factor model. The jointly estimated risk premiums $\lambda_{\Delta c}$, $\lambda_{\Delta\mu}$, and $\lambda_{\Delta\sigma}$ are, respectively, 0.57% (t = 4.69), 0.01% (t = 0.34), and −0.05% (t = −3.45), indicating that consumption growth and consumption volatility exposures have independently significant explanatory power. In our test (not shown here), we find that consumption growth highly correlates with shocks to expected consumption growth, with a correlation coefficient of around 0.57. The reason why exposures to expected consumption growth are not priced may be that their prices are absorbed in the risk prices of consumption growth. The signs of the risk premiums are consistent with investors' preference for early resolution of uncertainty ($\psi > 1/\gamma$). Since delta-hedged option portfolios are zero-cost portfolios, we further test whether the intercepts are zeros. Interestingly, we cannot reject that hypothesis, neither for the comprehensive model in specification IV nor for the less comprehensive models in specifications I, II, and III.

In the consumption-based model, the linear approximation of log changes in the wealth-consumption ratio suggests that the log-linearized SDF is approximately affine in consumption growth and the changes in its first two moments (see the Internet Appendix and Boguth and Kuehn (2013)). As an alternative to running FM regressions, we thus now use two-stage GMM to explicitly test the Euler equation of our model.[8] The second stage uses the optimal weighting matrix. The full-model moment condition can be written as:

$$E[(1 - b_{\Delta c}\Delta c_{t+1} - b_{\Delta\mu}\Delta\hat{\mu}_{t+1} - b_{\Delta\sigma}\Delta\hat{\sigma}_{t+1})R^i_{t+1}] = 0, \tag{15}$$

where $b_{\Delta c}$, $b_{\Delta\mu}$, and $b_{\Delta\sigma}$ are the SDF loadings. The GMM estimation of the Euler equation can generate useful results for elucidating the relation between the SDF and the consumption risk factors.

Table 4 Panel B presents the model estimates (both SDF loadings and implied risk premiums λ) and test statistics by using the 16 call option portfolios independently sorted on idiosyncratic underlying-stock volatility and options moneyness. Standard

[8] Iterative GMM yields results virtually identical to those reported.

Table 5
Predicting stock portfolio returns using consumption risk premiums.

	BM (L)				BM (H)
Average returns (%)					
Small	2.28	3.72	3.78	4.39	4.73
	2.86	3.65	3.97	4.11	4.31
	2.89	3.69	3.57	4.02	4.33
	3.18	3.13	3.39	3.86	3.95
Big	2.77	2.75	2.90	2.83	3.20
Market	2.84				
Model predicted returns (%)					
Small	6.80	6.84	6.89	6.70	7.57
	6.06	5.71	5.70	6.14	7.13
	5.84	5.38	5.50	5.51	5.66
	5.12	4.84	5.79	5.39	6.29
Big	4.91	4.60	4.38	5.62	4.90
Market	4.56				
ρ	0.54				

This table presents the average quarterly returns (Average Returns) of the market portfolio and the 25 value-weighted size-value stock portfolios, quarterly portfolio returns predicted by Eq. (14) (Model Predicted Returns), and the correlation coefficient between the average returns and model predicted returns (ρ). Returns are expressed in percentage. Monthly stock portfolio returns are downloaded from Professor Kenneth French's data library and are compounded into quarterly returns. Model predicted returns are calculated by multiplying consumption risk loadings of each portfolio and the corresponding consumption risk premiums estimated from option data plus the risk-free rate. The risk-free rate is the average quarterly risk-free rate over the sample period. The sample period is from July 1963 to December 2017.

errors are Newey–West (1987) adjusted with four lags. We find that, consistent with the FM regression results, the risk premiums for consumption growth and consumption volatility are significantly positive and negative, respectively. More specifically, the full model produces risk premium estimates for consumption growth, mean consumption growth, and consumption volatility of 0.21% (t = 5.11), −0.02% (t = −0.61), and −0.004% (t = −2.08), respectively. It further produces a mean absolute error (MAE) of 0.11% per quarter and an R^2 of over 99%, and the J-test of the over-identifying restrictions never rejects it ($p = 0.43$).

4.4. The pricing of consumption risks in stock and option returns

Because the consumption-based model is applicable to not only options but also to all assets in general, it is interesting to examine whether the consumption risk premiums estimated from options and presented in Table 4 also help to price other assets. To answer this question, we compare the mean returns of the market portfolio and the 25 size-value stock portfolios with the mean return predictions of our model and present the results in Table 5. We estimate each portfolio's exposures to consumption growth and the changes in mean growth and consumption volatility from a time-series regression of a portfolio's excess return on those three consumption moments. Combining the estimated exposures with the consumption risk premiums estimated from options, we calculate our consumption-based model's prediction for the mean return of each portfolio. For the market portfolio and the 25 size-value stock portfolios, we find that the correlation between the mean portfolio returns and the model predicted returns is 0.54. It is worth mentioning that the model predicted returns are generally larger than the average returns. It is well known in the literature that the magnitude of option returns is larger than that of stock returns and options embed larger risk premiums, which generate high model predicted returns. In our sample, the average Sharpe ratio of our option portfolios is −0.55, whereas the average Sharpe ratio of the 25 size-value stock portfolios is 0.20. The reward to risk ratio of option portfolios is thus more than two times as large as that of stock portfolios.

We further perform asset pricing tests using the joint cross-section of the 16 option portfolios and the 25 size-value stock portfolios. Table 6 presents the GMM model estimates (both SDF loadings and implied risk premiums λ) and test statistics. Standard errors are Newey–West (1987) adjusted with four lags. After adding stock portfolios to option portfolios, all three consumption risk premiums are significant. The full model generates risk premium estimates for consumption growth, mean consumption growth, and consumption volatility of 0.28% (t = 9.20), 0.04% (t = 4.04), and −0.01% (t = −2.33), respectively. It indicates that adding stock portfolios helps us to identify prices of shocks to expected consumption growth, and meanwhile the estimated risk premiums for the other two state variables are of similar magnitude with those estimated using only option portfolios in Table 4 Panel B.

Overall, our empirical analysis suggests that consumption risks can explain the negative relation between IVOL and the cross-section of delta-hedged option returns discovered by Cao and Han (2013). While these authors attribute the relation to market makers charging a premium for options that are difficult to delta hedge, the question remains why investors are content to pay the high premium. We show that investors are content to do so because options written on high IVOL stocks pay out more in adverse economic conditions, as evidenced by large negative consumption growth exposures and large positive consumption volatility exposures. Thus, these options are better hedging tools. Moreover, our analysis also suggests that consumption risks can explain the positive relation between delta-hedged call options and moneyness. Boyer and Vorkink (2014) argue that the positive relation arises because OTM options have more right-skewed payoffs than other options, which can be attractive to investors with non-standard preferences. Our evidence supports their finding that option moneyness is negatively related to option skewness. More importantly, we show that the

Table 6
Asset pricing tests with the joint cross-section of stock and option portfolios.

	Δc	$\Delta \mu$	$\Delta \sigma$	MAE	RMSE	J	R^2
b	1.81			0.42	0.54	29.07	96.18
	(7.08)					(0.90)	
λ	0.23						
	(7.08)						
b	1.90	4.01		0.41	0.52	29.39	96.43
	(7.45)	(2.72)				(0.87)	
λ	0.28	0.04					
	(9.28)	(5.71)					
b	1.78		−1.20	0.42	0.54	29.40	96.19
	(3.96)		(−0.82)			(0.87)	
λ	0.23		−0.00				
	(4.48)		(−1.27)				
b	1.84	4.53	−3.17	0.40	0.52	30.55	96.48
	(8.52)	(3.13)	(−1.95)			(0.80)	
λ	0.28	0.04	−0.01				
	(9.20)	(4.04)	(−2.33)				

This table reports the GMM estimates of the moment conditions in Eq. (15) using the 16 quarterly IVOL-moneyness call option portfolio returns combined with the 25 value-weighted size-value stock portfolio returns. We report both the b estimates as well as the implied risk premiums (λ). MAE and RMSE (in percentage) refer to the mean absolute pricing error and the root mean squared error, respectively. Newey–West (Newey and West, 1987) adjusted t-statistics using four lags are reported in parentheses, and p-values for J-statistics are shown in parentheses below the associated J-statistics. The sample period is from January 1996 to December 2017.

low returns of low-moneyness and/or high-skewness options can be explained by these options being better instruments to hedge against adverse economic conditions. Here, adverse economic conditions not only include shocks to consumption growth but also variations in expected consumption growth and consumption volatility that can lead to a rise in the SDF under recursive utility. Moreover, we find that consumption risk premiums estimated from options data do surprisingly well in pricing stock portfolios and adding stock portfolios in option portfolios further helps us to identify prices of shocks to expected consumption growth.

5. Conclusion

We study the impact of consumption risks on the cross-section of delta-hedged option returns using portfolios sorted on IVOL and options moneyness. Our consumption-based pricing factors consist of consumption growth, an estimate of its conditional expectation, and an estimate of its conditional volatility. The three factors explain the cross-section of delta-hedged option returns well and support a risk-based explanation for option returns. Consumption growth commands a positive risk premium, whereas consumption volatility commands a negative risk premium. Our evidence suggests that, in a representative-agent economy with recursive preferences, the agent prefers early resolution of uncertainty. We further suggest that options written on high IVOL stocks and options with a low moneyness earn more negative returns than other options because they provide a better hedge against bad macroeconomic conditions. Moreover, consumption risk premiums estimated from option portfolios can predict stock portfolio returns well, and adding stock portfolios in option portfolios further identifies a significantly positive expected consumption growth risk premium. Overall, our empirical findings are robust to the choice of test assets, option maturities, and the testing horizon. Taken together, our analysis provides a foundation for consumption risks explaining the cross-section of delta-hedged option returns.

CRediT authorship contribution statement

Shuwen Yang: Software, Validation, Formal analysis, Investigation, Data curation, Visualization, Writing – original draft. **Kevin Aretz:** Validation, Formal analysis, Supervision, Writing – review & editing. **Hening Liu:** Conceptualization, Methodology, Software, Project administration, Supervision, Writing – review & editing. **Yuzhao Zhang:** Validation, Formal analysis, Supervision, Writing – original draft.

Acknowledgment

The funder to acknowledge is China Scholarship Council, and the grant ID is 201708060485.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jempfin.2022.10.001.

References

Abel, A.B., 1999. Risk premia and term premia in general equilibrium. J. Monetary Econ. 43, 3–33.
Aretz, K., Gazi, A., 2021. The Early Exercise Risk Premium. Working Paper.
Aretz, K., Lin, M.T., Poon, S.H., 2022. Moneyness, underlying asset volatility, and the cross-section of option returns. Review of Finance forthcoming.
Bakshi, G., Kapadia, N., 2003. Delta-hedged gains and the negative market volatility risk premium. Rev. Financ. Stud. 16, 527–566.
Bansal, R., Yaron, A., 2004. Risks for the long run: A potential resolution of asset pricing puzzles. J. Finance 59, 1481–1509.
Boguth, O., Kuehn, L.A., 2013. Consumption volatility risk. J. Finance 68, 2589–2615.
Boyer, B.H., Vorkink, K., 2014. Stock options as lotteries. J. Finance 69, 1485–1527.
Breeden, D.T., 1979. An intertemporal asset pricing model with stochastic consumption and investment opportunities. J. Financ. Econ. 7, 265–296.
Cao, J., Han, B., 2013. Cross section of option returns and idiosyncratic stock volatility. J. Financ. Econ. 108, 231–249.
Coval, J.D., Shumway, T., 2001. Expected option returns. J. Finance 56, 983–1009.
Dew-Becker, I., Giglio, S., Le, A., Rodriguez, M., 2017. The price of variance risk. J. Financ. Econ. 123, 225–250.
Drechsler, I., Yaron, A., 2011. What's vol got to do with it. Rev. Financ. Stud. 24, 1–45.
Epstein, L.G., Zin, S.E., 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. Econometrica 57, 937–969.
Fama, E.F., MacBeth, J.D., 1973. Risk, return, and equilibrium: Empirical tests. J. Polit. Econ. 81, 607–636.
Gallant, A.R., Jahan-Parvar, M.R., Liu, H., 2019. Does smooth ambiguity matter for asset pricing? Rev. Financ. Stud. 32, 3617–3666.
Goyal, A., Saretto, A., 2009. Cross-section of option returns and volatility. J. Financ. Econ. 94, 310–326.
Hamilton, J.D., 1994. Time Series Analysis. Princeton University Press, Princeton, NJ.
Hansen, L.P., 1982. Large sample properties of generalized method of moments estimators. Econometrica 50, 1029–1054.
Hu, G., Jacobs, K., 2020. Volatility and expected option returns. J. Financ. Quant. Anal. 55, 1025–1060.
Jagannathan, R., Wang, Y., 2007. Lazy investors, discretionary consumption, and the cross-section of stock returns. J. Finance 62, 1623–1661.
Kim, C.J., Nelson, C.R., 1999. Has the U.S. economy become more stable? A Bayesian approach based on a markov-switching model of the business cycle. Rev. Econ. Stat. 81, 608–616.
Lettau, M., Ludvigson, S., 2001. Resurrecting the (C)CAPM: A cross-sectional test when risk premia are time-varying. J. Polit. Econ. 109, 1238–1287.
Lettau, M., Ludvigson, S.C., Wachter, J.A., 2008. The declining equity premium: What role does macroeconomic risk play? Rev. Financ. Stud. 21, 1653–1687.
Liu, H., Zhang, Y., 2022. Financial uncertainty with ambiguity and learning. Manage. Sci. 68, 2120–2140.
Mankiw, N.G., Shapiro, M.D., 1986. Risk and return: Consumption beta versus market beta. Rev. Econ. Stat. 68, 452–459.
McConnell, M.M., Perez-Quiros, G., 2000. Output fluctuations in the United States: What has changed since the early 1980's? Amer. Econ. Rev. 90, 1464–1476.
Newey, W.K., West, K.D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703–708.
Parker, J.A., Julliard, C., 2005. Consumption risk and the cross section of expected returns. J. Polit. Econ. 113, 185–222.
Romeo, T., 2015. Consumption volatility and the cross-section of stock returns. Rev. Finance 19, 367–405.
Veronesi, P., 1999. Stock market overreaction to bad news in good times: A rational expectations equilibrium model. Rev. Financ. Stud. 12, 975–1007.
Yogo, M., 2006. A consumption-based explanation of expected stock returns. J. Finance 61, 539–580.