Journal of Economic Dynamics & Control 34 (2010) 2578–2600

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

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

Equilibrium open interest ☆

Kenneth L. Judd [a], Dietmar P.J. Leisen [b,*]

[a] Stanford University, Hoover Institution and NBER, Stanford, CA 94305, United States
[b] University of Mainz, Faculty of Law and Economics, 55099 Mainz, Germany

ARTICLE INFO

Article history:
Received 15 June 2009
Accepted 9 July 2010
Available online 4 August 2010

JEL classification:
G11
G12
G13

Keywords:
Option demand
Open interest
Co-skewness
Skewness preference

ABSTRACT

This paper analyses what determines an individual investor's risk-sharing demand for options and, aggregating across investors, what the equilibrium demand for options. We find that agents trade options to achieve their desired skewness; specifically, we find that portfolio holdings boil down to a three-fund separation theorem that includes a so-called skewness portfolio that agents like to attain. Our analysis indicates also, however, that the common risk-sharing setup used for option demand and pricing is incompatible with a stylized fact about open interest across strikes.

1. Introduction

Option volume and open interest have steadily increased over the last decades and reached tremendous amounts.[1] However, the option literature focuses on the pricing of options under risk-sharing assumptions but largely ignores the demand in such contracts. This paper determines portfolio holdings in options and analyses what drives risk-sharing option demand.

In this paper we study a series of single-period, partial and general equilibrium, exchange economies. Each economy consists of two agents that can trade a stock and a portfolio of call options with various strikes written on the stock. While the stock is in unit supply, the options are in zero net-supply. To determine demand we use the small-noise expansion technique introduced by Samuelson (1970) and recently extended by Judd and Guu (2001). It provides closed-form approximations for demand and prices that capture their salient features.

In partial equilibrium we show that a security's risk-premium is proportional to its covariance with the agent's portfolio payoff and her squared portfolio payoff. The proportionality constants are functions of agent's risk-preferences which we determine. In general equilibrium we introduce a new portfolio, that is as close as possible to the squared stock payoff, the so-called skewness portfolio. Agents' portfolio holdings boil down to a three-fund separation theorem: agents trade the bond, the stock (the surrogate for the market portfolio) and the skewness portfolio (the surrogate for the squared market

☆ We are grateful for comments and suggestions from seminar participants at Stanford University and McGill University, and from conference participants at the North American Winter Meetings of the Econometric Society 2003 and at the European Summer Symposium in Financial Markets 2003. We are particularly grateful to Jun Pan and Raman Uppal. All remaining errors are ours.

* Corresponding author. Tel.: +49 6131 3925542; fax: +49 6131 3929905.
E-mail addresses: judd@Hoover.Stanford.EDU, leisen@Uni-Mainz.DE (D.P. Leisen).

[1] Open interest in a financial security denotes the total number of that contract held long; it is quoted at the end of each day for all financial contracts traded. According to the Bank for International Settlements, the options turnover in 2007 was 4525.3 million contracts with an outstanding notional principle of 52,521.3 billion US dollars at the end of 2007. To put this into perspective, the International Monetary Fund reports that world gross domestic product in 2007 was 54,311.6 billion US dollars.

0165-1889/$ - see front matter
doi:10.1016/j.jedc.2010.07.006

portfolio); risk-premiums exhibit a beta-gamma representation where idiosyncratic market and squared market risk are not priced. Finally we confront our results to a stylized fact about open interest and, surprisingly, find that a risk-sharing approach with the usual risk-preferences is incompatible with it.

The importance of contingent claims for risk-sharing[2] is well known since Arrow (1953) and Debreu (1959). Ross (1976) pointed out that options complete the market: this suggests that agents demand options. The option literature therefore often ignores option demand and simply focuses on their pricing: the workhorse is the representative agent who prices assets based on the marginal utility they provide but cannot justify demand in zero net-supply securities. This paper sticks to the risk-sharing rationale for options, but instead of starting at the aggregation level, we start at the level of individual agents.

Our paper makes four contributions. First, in a partial equilibrium, we explain how buy-and-hold investors will position themselves in a stock and a portfolio of options. Using the small-noise expansion technique we link demand to closed-form expressions in mean, variance and skewness of payoff distributions. Our analysis only requires that investors know the first three (co-)moments of securities' payoffs, but does not assume either complete or incomplete markets, risk-preferences other than expected utility or a type of distributions. Previous studies, however, have relied on specific assumptions: agents have CPRA preferences; prices follow (jump-) diffusions; often a single option is studied. In a diffusion model with CPRA investors Garleanu et al. (2009) find that demand pressures increase option prices by an amount proportional to the variance of the unhedgeable part. In a jump-diffusion with exogenously specified pricing kernel and one option that completes the market, Liu and Pan (2003) find intuitive closed-form expressions for demand by CPRA investors; they document empirically that options improve the investor's risk-return profile. Using the nonparametric approach of Brandt (1999) and Ait-Sahalia and Brandt (2001) to determine holdings in a single option, Driessen and Maenhout (2007) document empirically that CPRA investors will always short some types of options.

Second, we analyze in detail the aggregation across two agents to get insights into the reasons for option demand. Aggregation leads to tedious mathematics, see Detemple and Selden (1991), and, so, the literature has focused on special cases and on the question which agents are buying (selling) options, see Leland (1980), Brennan and Solanki (1981), Franke et al. (1998), Dieckmann and Gallmeyer (2005), Bates (2008), and Weinbaum (2009).[3] Using the small-noise expansion technique we can capture the salient features of demand in closed-form expressions. We provide the following explanation for risk-sharing option demand: If agents care about asymmetric events like liquidity shocks and market crashes, such preferences show up as preferences over portfolio skewness; agents demand call options because they alter the skewness of their portfolio payoff by trading the upper tail of the stock distribution. This shows up in their holdings of the so-called skewness portfolio.

Third, as we determine the equilibrium allocation, we also derive risk-sharing equilibrium prices; therefore, our paper adds to the option pricing literature. Based on the important contribution of Harvey and Siddique (2000) the literature has renewed interest in skewness for securities pricing. Our paper finds that co-skewness is relevant for option pricing, in line with Vanden (2006) and Bakshi et al. (2001).

Finally, we contribute to our understanding about the driving forces on option markets. Our theoretical risk-sharing analysis points to an incompatibility of risk-sharing with the stylized fact about the open interest curve across strikes. This indirectly adds support to other rationales for option demand, i.e. information trading and differences of beliefs. Our paper fits into the recent literature that looks at information advantages that manifest themselves in option trade: For example, Pan and Poteshman (2006) document that option volume may predict future stock price movements; Buraschi and Jiltsov (2006) find that heterogeneous beliefs can help explain option volume.

The remainder of the paper is as follows: Section 2 describes our setup, while the following section discusses its features and our small-noise expansion approach. Sections 4 and 5 discuss separately the incomplete and complete markets cases. Section 6 discusses the stylized fact about the shape of open interest across strikes and confronts it with the risk-sharing implications from incomplete and complete markets. Section 7 discusses the robustness of our approach, reviewing our assumptions and resulting errors. Section 8 concludes the paper. We postpone mathematical details to the appendix.

2. The model

We study a sequence of financial economies, parameterized by the parameter $0 \le \sigma \le 1$. Each $\sigma-$economy has two dates, zero and one, and is populated by two agents a=1,2 that can invest at date 0 in a riskfree bond, a stock and call options with maturity at date 1, written on the stock.

The total supply of stock is one unit, which is infinitely divisible; all call options are in zero net-supply. For simplicity we set the riskfree interest rate to 0, and assume the bond supply is fully elastic. The date 1 payoff from the stock in the $\sigma-$economy is

$$Y_S(\sigma) = 1 + \sigma\varepsilon_0. \tag{1}$$

[2] The literature has also argued that options are traded based on differential information, see, e.g. John et al. (1992), Back (1993), Brennan and Cao (1996), Easley et al. (1998). Option demand has also been motivated based on differences of opinion, see, e.g. Kraus and Smith (1996), and Carr and Madan (2001), and based on frictions trading the underlying security, see, e.g. Haugh and Lo (2001).

[3] There is also a large literature that looks at feedback effects between option markets, stock markets, interest rates and volatility, including Bhamra and Uppal (2009), Kogan and Uppal (2001) and Benninga and Mayshar (2000).

Here ε_0 is a zero-mean random variable that is exogenous to our analysis; its distribution is known to both agents, only takes values in the interval[4] $(-1,1)$ and its density has continuous support.

Call options written on the stock have date 1 payoff

$$Y_K(\sigma) = (Y_S(\sigma) - K(\sigma))^+, \tag{2}$$

they are fully characterized by the parameter K which we assume $0 < K < 2$. The strike of option K in the σ-economy is

$$K(\sigma) = 1 + \sigma \cdot (K-1), \tag{3}$$

i.e. it is a function of the parameter K. Throughout the paper it will be clear from the context whether we use the strike $K(\sigma)$ or its (characterizing) parameter K. Note that we can write simply

$$Y_K(\sigma) = \sigma \cdot (1 + \varepsilon_0 - K)^+. \tag{4}$$

We denote the date 0 price of the stock and all options K by $P_S(\sigma)$ and $P_K(\sigma)$, respectively, and decompose

$$P_S(\sigma) = E[Y_S(\sigma)] - \pi_S(\sigma) \cdot \sigma^2 \quad \text{for the stock,} \tag{5}$$

$$P_K(\sigma) = E[Y_K(\sigma)] - \pi_K(\sigma) \cdot \sigma^2 \quad \text{for option } K. \tag{6}$$

The terms $\pi_S(\sigma)\sigma^2, \pi_K(\sigma)\sigma^2$ are the difference between securities' expected payoff and their price, i.e. they describe the *risk-premiums* because we set the interest rate to zero.

Each agent is endowed with one-half units of the stock in all σ-economies. We assume, agents have preferences over date 1 wealth $W_{1a}(\sigma)$, only, and these can be represented by von Neumann–Morgenstern utility functions $E[u_a(W_{1a}(\sigma))]$, where u_a is an increasing, strictly concave and sufficiently often differentiable function for both agents $a=1,2$.

We denote $d_{aS}(\sigma)$ and $d_{aK}(\sigma)$, respectively, the number of units held of the stock and of options K. Since the bond supply is fully elastic, agents use the bond demand to equate today's portfolio value with their initial wealth $W_{0a}(\sigma) = P_S(\sigma)/2$. The agent then chooses the strategy that maximizes her expected utility $E[u_a(W_{1a}(\sigma))]$ over all holdings in the risky securities. In our *partial* equilibrium analysis we look at how agents set demand in response to the market risk-premiums (prices) they face.

We are mainly interested in market-clearing demand and prices. Both agents $a=1,2$ differ in their preferences (utility functions u_a) and trade competitively the available risky securities. The *general* equilibrium in the σ-economy consists of prices $P_S(\sigma), P_K(\sigma)$ for the available risky securities, and portfolio demand $d_{aS}(\sigma), d_{aK}(\sigma)$ for both agents $a=1,2$, that maximize agent's expected utility, and for which the stock and option markets for all K clear, i.e. $d_{1S}(\sigma) + d_{2S}(\sigma) = 1$ and $d_{1K}(\sigma) + d_{2K}(\sigma) = 0$ for all K.

3. Expansion technique

This section first discusses the main features of our sequential modeling and then explains that it is a sequence of classical risk-sharing economies. Finally we discuss how the sequential features will be used to approximate allocations in closed-form expressions.

3.1. Model features

The expansion is set up such as to keep the cumulative probability of option exercise unaffected by changes in $\sigma > 0$, i.e. $\text{Prob}[Y_S(\sigma) \geq K(\sigma)] = \text{Prob}[1 + \varepsilon_0\sigma \geq 1 + \sigma(K-1)] = \text{Prob}[\varepsilon_0 \geq K-1]$.

Fig. 1 illustrates our expansion for three strikes $K_1(\sigma) \leq K_2(\sigma) \leq K_3(\sigma)$ set at the quartiles of the stock distribution. When $\sigma = 0$, the support of the distribution of $Y_S(0)$ is concentrated at 1 and all strikes collapse to 1, so that none of them pays off. As σ increases, the support of $Y_S(\sigma)$ increases linearly:[5] For any $\sigma > 0$ the support is $(1-\sigma, 1+\sigma)$; this is shown through the bold lines in Fig. 1. Furthermore, as σ increases, the location of the quartiles increases linearly around 1 and to keep strikes at the quartiles we expand strikes linearly, see Eq. (3); this is shown through the dashed lines in Fig. 1.

A characteristic feature of our model for payoffs, prices and risk-premiums is that $\sigma = 0$ corresponds to an economy without risk. Then, today's price and tomorrow's payoff of the stock (option K) will be 1 and 1 (0 and 0), respectively; this means that the return on the stock in this "no-risk" economy will be identical to the bond return and options vanish entirely.

For any strictly positive $\sigma > 0$, the payoff from the stock and options is risky. Note that the variance of the stock and option K is $var(\varepsilon_0)\sigma^2$ and $var((1+\varepsilon_0-K)^+)\sigma^2$, respectively. Therefore, an increase in σ will increase the variance of the stock and option payoff proportionally in σ^2. If $\pi_S(\sigma), \pi_K(\sigma)$ would be constant in σ, then $\pi_S(\sigma)\sigma^2, \pi_K(\sigma)\sigma^2$ would give risk

[4] For technical reasons the support of ε_0 must be bounded; to simplify our presentation we assume it only takes values in the interval $(-1,1)$. Being bounded also ensures that stock payoffs are non-negative.

[5] It is *not* feasible to keep strikes fixed since we are interested in the case where $\sigma \to 0$. If we would keep strikes fixed, then for sufficiently small σ the probability of exercise would be zero and all options would vanish, because the support of ε_0 is bounded.

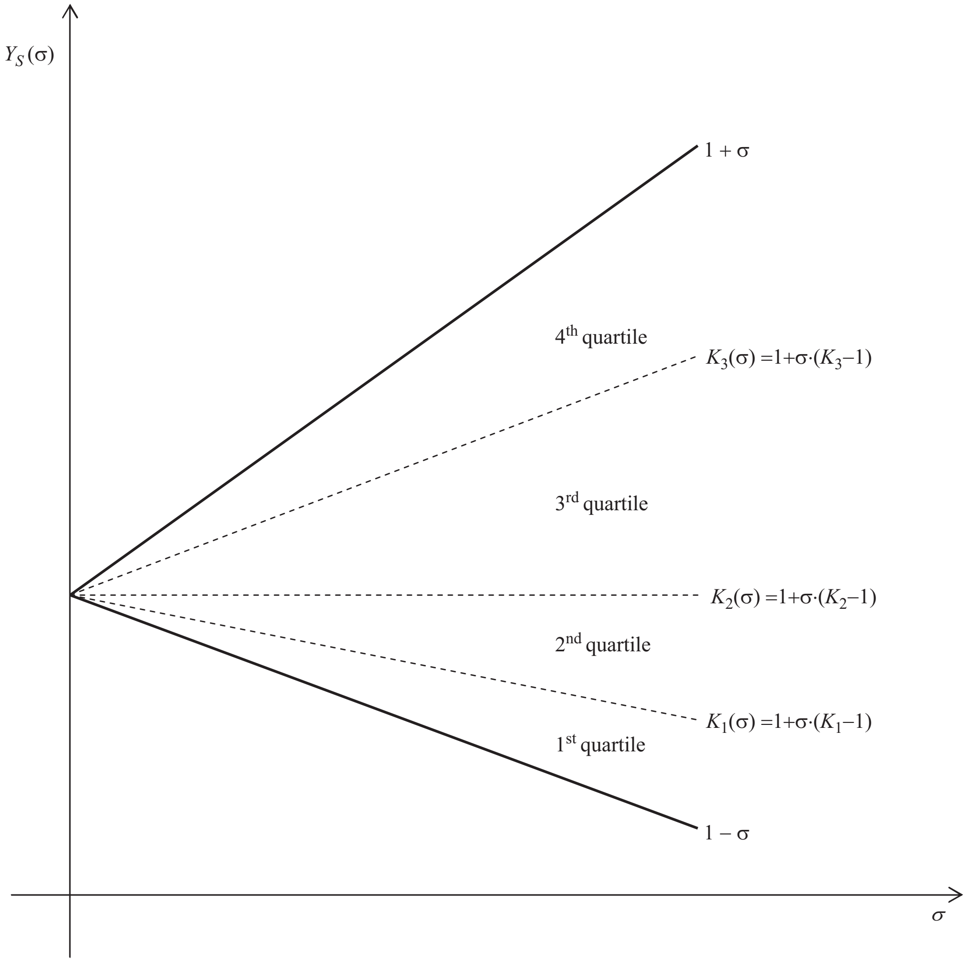

Fig. 1. Expanding the strikes such that the probability of exercise remains unaffected: strikes $K_1(\sigma),K_2(\sigma),K_3(\sigma)$ are at the quartiles in the σ−economy.

premiums that are proportional to variance. We allow $\pi_S(\sigma),\pi_K(\sigma)$ to depend on σ to capture dependence on higher-order moments of securities' distribution.

At first, it might appear odd to look at a sequence of economies. Note, however, that for any given $0 \le \sigma \le 1$, we model stock payoffs in a classical way and the option payoff is just a standard call option payoff. Furthermore, in each σ−economy, risk-sharing determines prices and demand. Therefore, our setup is just a sequence of standard risk-sharing economies used to determine equilibrium demand and prices.

3.2. Approach

To analyze demand, it is necessary to start at the level of individual agents. However, such analyses rarely lead to closed-form expressions; in general, they are cumbersome to carry out and results are hard to interpret. Following Judd and Guu (2001) we expand stock demand, option demand and their risk-premiums into Taylor-series around $\sigma = 0$, i.e.

$$d_{aS}(\sigma) = d_{aS}(0) + \frac{\partial d_{aS}}{\partial \sigma}(0)\cdot\sigma + \frac{\partial^2 d_{aS}}{\partial \sigma^2}(0)\cdot\sigma + \cdots, \tag{7}$$

$$\pi_S(\sigma) = \pi_S(0) + \frac{\partial \pi_S}{\partial \sigma}(0)\cdot\sigma + \frac{\partial^2 \pi_S}{\partial \sigma^2}(0)\cdot\sigma + \cdots, \tag{8}$$

and similarly for any option with strike K. We will see throughout this paper that these expansion terms *can* be determined in closed-form. Of course, nothing is gained to previous studies, if we do calculate all these expansion terms (an infinite number). But this is really not necessary to capture the salient features of option demand in closed-form. Instead, we look at zero and first order demand approximations

$$d_{aS}^{zero}(\sigma)=d_{aS}(0) \quad \text{and} \quad d_{aS}^{first}(\sigma)=d_{aS}(0)+\frac{\partial d_{aS}}{\partial \sigma}(0)\sigma,$$

and similarly for any option with strike K. (We sometimes refer to $d_{aS}(0)$, $d_{aK}(0)$ as zero order demand expansion terms, and to $\partial d_{aS}/\partial\sigma(0),\partial d_{aK}/\partial\sigma(0)$ as first order demand expansion terms. Notation is similar for the risk-premiums.) These terms will permit us to capture the salient features of option demand.

We refer to the expansion as the *small-noise expansion* technique; it was first introduced by Samuelson (1970) to motivate the use of mean–variance analysis by studying agent's portfolio allocations in the limit $\sigma\rightarrow 0$. While Samuelson's analysis focused on the zero-order term, the technique used here is an extension developed by Judd and Guu (2001) that is asymptotically valid for all expansion terms.

Throughout our analysis we do *not* rely on specific distributions to gain insights into the structure of demand and risk-premiums. Almost never our analysis will therefore require modeling distributions. Yet, at times we want to illustrate demand and for this we need to adopt specific distributions ε_0 for the stock. We will then look at three distributions truncated to the interval $(-1,1)$: the uniform distribution, the normal distribution (mean 0, variance 1/3), and the (transformed) lognormal $\exp(X)-1$. For the last, X is a random variable with mean 0.0036 and variance 1/3, truncated to the interval $(-\infty,\ln 2)$. (The random variable $\exp(X)-1$ then has mean 0, variance approximately equal to 1/3 and is defined on the interval $(-1,1)$.)

The uniform distribution on the interval $(-1,1)$ has exactly variance 1/3; the truncated normal and lognormal distributions have approximately the same variance 1/3. For $\sigma=1$, this is the payoff variance and so its volatility is approximately 57.7%. This is within the range of typical stock volatility. Because volatility scales linearly in time-length and σ has the same property, we will use σ to scale the distributions to the (short) time-length of one week to one month we have in mind: a typical value for σ will then be from 1/52 to 1/12.

4. Incomplete markets

This section looks at sequences of incomplete financial markets, i.e. only a finite number N of options are traded with parameters $0<K_1<\cdots<K_N<2$. Because we make extended use of vector/matrix notation throughout this section, we refer to the stock as security number 0, to the option with parameter K_j ($j=1,\ldots,N$) as option j, and to their payoffs, prices, risk-premiums and demand as $Y_j(\sigma),P_j(\sigma),\pi_j(\sigma)\sigma^2$, and $d_{aj}(\sigma)$, respectively ($j=0,\ldots,N$).

4.1. Notation

Unless necessary to prevent confusion we drop the σ notation to simplify our exposition. We denote by Y the $(N+1)$ dimensional payoff vector for the risky securities, by V its $(N+1)\times(N+1)$ dimensional *covariance matrix*, by χ the $N+1$ dimensional *co-skewness vector* and for given demand vector $\tilde{d}$ by $\zeta(\tilde{d})$ the $N+1$ dimensional *(third order) co-moment vector*, i.e. we set for $j,k=0,\ldots,N$:

$$V_{jk}=Cov(Y_j,Y_k), \tag{9}$$

$$\chi_j=Cov(Y_j,((Y_0-E[Y_0])^2) \quad \text{and} \tag{10}$$

$$\zeta_j(\tilde{d})=Cov(Y_j,(\tilde{d}^{\perp}\cdot(Y-E[Y]))^2). \tag{11}$$

Note that the variance–covariance matrix V of the $N+1$ securities is not singular for $\sigma>0$, i.e. V is invertible, since the support of ε_0 is continuous. The (third order) co-moment vector $\zeta(\tilde{d})$ describes how the squared (de-meaned) payoff from a portfolio $\tilde{d}$ covaries with each security; the co-skewness vector χ describes how the squared (de-meaned) stock payoff covaries with each security. $\zeta(\tilde{d})$ will be used in the partial equilibrium of Section 4.2 only; in the general equilibrium of Section 4.3 co-skewness χ will take its role for prices and demand.

We define the *market portfolio* as the portfolio consisting of one unit of stock and no options; it is defined by the vector X_M of holdings with $X_{M0}=1$ and $X_{Mj}=0$ for $j=1,\ldots,N$. Note that the stock is the only security that is not in zero net-supply. In our setup, the stock plays the role of the market portfolio, and, therefore, we think of the stock as the "market portfolio" for illustrative purposes. We also introduce a second portfolio which we call the *skewness portfolio*; it is defined by the vector of holdings $X_S=V^{-1}\chi$.

Finally, we define for each agent $a=1,2$ the *risk-tolerance* τ_a and the *skew-tolerance* ρ_a by

$$\tau_a=-\frac{\dfrac{\partial u_a}{\partial w}\left(\dfrac{1}{2}\right)}{\dfrac{\partial^2 u_a}{\partial w^2}\left(\dfrac{1}{2}\right)}, \quad \rho_a=\frac{\tau_a^2}{2}\frac{\dfrac{\partial^3 u_a}{\partial w^3}\left(\dfrac{1}{2}\right)}{\dfrac{\partial u_a}{\partial w}\left(\dfrac{1}{2}\right)}. \tag{12}$$

These two terms are ratios of derivatives of agent's utility functions; they are evaluated at one-half, so, these two preference parameters will be treated as (non-random) numbers. (The value one-half stands for the endowment value in the "no-risk" economy $\sigma=0$.) The risk-tolerance τ_a consists of first- and second-order derivatives of agents' utility functions and is an usual term in the financial economics literature. The skew-tolerance adds a third-order derivative of agents' utility functions and appeared first in Judd and Guu (2001); it is related to the prudence concept of Kimball (1990).

4.2. Partial equilibrium

Appendix A.3 derives the individual asset allocation across options and the stock for given risk-premiums:

Theorem 1. *Zero and first order demand approximations for agent i are*

$$d_a^{zero}(\sigma)=\tau_a\cdot V^{-1}\pi^{zero}(\sigma)\sigma^2 \tag{13}$$

and

$$d_a^{first}(\sigma)=\tau_a\cdot V^{-1}\pi^{first}(\sigma)\sigma^2+\frac{\rho_a}{\tau_a}\cdot V^{-1}\zeta(d_a^{zero}(\sigma)). \tag{14}$$

The zero order demand approximation is exactly the well-known one for a mean–variance agent, see, e.g. Huang and Litzenberger (1988). In this sense, this term describes mean–variance demand. Indeed, Samuelson (1970) originated the small-noise expansion to motivate the use of mean–variance analysis as an approximation for complex portfolio choice situations.

As the agent has preferences over skewness, in general, a correction to the mean–variance based demand term is necessary. The term $\zeta(d_a^{zero}(\sigma))$ in the first order demand approximation describes how each security covaries with the squared portfolio payoff coming from the zero order demand $d_a^{zero}(\sigma)$. This is a third order cross-moment that captures (co-) skewness across all securities. Therefore, the first order demand approximation captures agents' tradeoffs between securities risk-premiums, variances V and third-order (co-)moments $\zeta(d_a^{zero}(\sigma))$. Here, the ratio of skew- to risk-tolerance ρ_a/τ_a characterizes the marginal rate of substitution between variance and skewness risks. Overall, our analysis shows that our first order demand approximation provides a mean–variance–skewness theory of demand.

Rewriting this, the resulting first-order approximation for the risk-premium $\pi^{first}(\sigma)\sigma^2$ that induces a given zero and first order demand[6] by agent a is

$$\pi^{first}(\sigma)\sigma^2=\frac{1}{\tau_a}V\cdot d_a^{first}(\sigma)+\frac{\rho_a}{\tau_a^2}\cdot\zeta(d_a^{zero}(\sigma)). \tag{15}$$

If we ignore for a moment the difference between zero and first order demand approximations and denote both by $\tilde{d}_a$, then we get using the portfolio payoff $Y_a=\tilde{d}_a^{\perp}\cdot Y$

$$\pi_j^{first}(\sigma)\sigma^2=\frac{1}{\tau_a}Cov(Y_j,Y_a)+\frac{\rho_a}{\tau_a^2}\cdot Cov(Y_j,Y_a^2),$$

i.e. how a security's payoff covaries with the individual portfolio payoff and the squared individual portfolio payoff determines its first order risk-premium.

4.3. General equilibrium

In equilibrium, in each σ–economy, agents have to agree on market-clearing prices, i.e. $d_{10}(\sigma)+d_{20}(\sigma)=1$ for the stock and $d_{1j}(\sigma)+d_{2j}(\sigma)=0$ for option j=1,…,N. For the zero and first order expansion terms of $d_a(\sigma)$ this translates into the conditions $d_{10}(0)+d_{20}(0)=1$ for the stock, $d_{1j}(0)+d_{2j}(0)=0$ for option j=1,…,N, and $\partial d_{1j}/\partial\sigma(0)+\partial d_{2j}/\partial\sigma(0)=0$ for (all) securities j=0,…, N. Appendix A.4 derives the following characterization.

Theorem 2 (*First order equilibrium allocation*). *The first order stock holdings are*

$$d_{10}^{first}(\sigma)=\frac{\tau_1}{\tau_1+\tau_2}-\tau_1\tau_2\frac{\rho_1-\rho_2}{(\tau_1+\tau_2)^3}(V^{-1}\chi)_0,\quad\text{for the first agent,}$$

$$d_{20}^{zero}(\sigma)=\frac{\tau_2}{\tau_1+\tau_2}+\tau_1\tau_2\frac{\rho_1-\rho_2}{(\tau_1+\tau_2)^3}(V^{-1}\chi)_0,\quad\text{for the second agent,}$$

the first-order holdings in option j=1,…, *N are*

$$d_{1j}^{first}(\sigma)=-\tau_1\tau_2\frac{\rho_1-\rho_2}{(\tau_1+\tau_2)^3}(V^{-1}\chi)_j,\quad\text{for the first agent,}$$

$$d_{2j}^{first}(\sigma)=\tau_1\tau_2\frac{\rho_1-\rho_2}{(\tau_1+\tau_2)^3}(V^{-1}\chi)_j,\quad\text{for the second agent,}$$

[6] Demand pressures by one agent on the prices set by another agent have been looked at by Garleanu et al. (2009) and linked to the unhedgeable part of each option. In our setup, the hedge ratio is not an intrinsic term and so we refrain from linking our results directly to theirs.

and the first order risk-premiums are

$$\pi_j^{\text{first}}(\sigma)\sigma^2 = \frac{1}{\tau_1+\tau_2}V_{0j} - \frac{\tau_1\rho_1+\tau_2\rho_2}{(\tau_1+\tau_2)^3}\chi_j \quad \text{for securities } j=0,\ldots,N.$$

We saw earlier that the zero order demand captures mean–variance demand and that the first-order demand adds skewness to the analysis. The results of Theorem 2 should then come at no surprise: The zero order risk-premium corresponds to the CAPM pricing term and describes the mean–variance part of the risk-premium. The first-order term depends on co-skewness χ_j, i.e. on the covariance of security j's payoff with the squared (de-meaned) stock payoff; therefore we call the first-order term the (co-)skewness part of the risk-premium.

Note that the zero order demand term vanishes for options. This result is in line with mean–variance analyses: there, options are not traded because optimal spanning can be achieved by trading the stock and the bond. Put differently, a simple mean–variance analysis is not sufficient to analyze option demand. For this purpose (co-)skewness needs to be analyzed, i.e. a first-order analysis is necessary. In fact we see that options will be traded as long as agents differ in their skew-tolerances $\rho_1 \neq \rho_2$. To understand demand further, we write demand in the form of the following three-fund separation.

Theorem 3 (*Three-fund separation*). *In the first-order approximation of equilibrium demand agents hold three funds*: *the riskfree bond*, *and the market and skewness portfolios. The first* (*second*) *agent holds* $\tau_1/(\tau_1+\tau_2)$ $(\tau_2/(\tau_1+\tau_2))$ *units of the market portfolio and* $(-\tau_1\tau_2(\rho_1-\rho_2)/(\tau_1+\tau_2)^3)(\tau_1\tau_2(\rho_1-\rho_2)/(\tau_1+\tau_2)^3)$ *units of the skewness portfolio.*

Note that holdings of the market portfolio are zero-order terms in σ, while holdings of the skewness portfolio are first-order demand terms.[7] The stock is the only risky security that is in positive supply. Therefore we say it captures the systematic "market" risk and the squared stock the systematic squared market risk component. In line with our three-fund separation we find that only systematic risks are priced.

Theorem 4 (*Systematic risk pricing*). *Using the beta-vector* $\beta_j = V_{0j}/V_{00}$ *and the gamma-vector* $\gamma_j = \chi_j/\chi_0$, *the first-order approximation of the risk-premium* $\pi(\sigma)\sigma^2$ *is characterized by*

$$\pi^{first}(\sigma)\sigma^2 = \beta\frac{V_{00}}{\tau_1+\tau_2} - \gamma\cdot\frac{\tau_1\rho_1+\tau_2\rho_2}{(\tau_1+\tau_2)^3}\chi_0. \tag{16}$$

Overall, the beta-gamma representation shows that the risk-premium is based on covariances with the market risk and with the squared market risk. This theoretical result points to the importance of (co-)skewness risk for pricing and confirms earlier results by Huang and Litzenberger (1988, p. 100). It is also in line with the empirical evidence, see Kraus and Litzenberger (1976), and more recently Harvey and Siddique (2000), Dittmar (2002), Chang et al. (2006) and Vanden (2006).

Eq. (6) together with the representation of the option risk-premium in Theorem 2 implies the following first-order approximation for the price of option $j=1,\ldots,N$:

$$E[Y_j(\sigma)] - \frac{1}{\tau_1+\tau_2}V_{0j} + \frac{\tau_1\rho_1+\tau_2\rho_2}{(\tau_1+\tau_2)^3}\chi_j. \tag{17}$$

To illustrate the pricing implications we look at the so-called implied volatility for which this price matches the option price that can be calculated from the Black–Scholes formula. We perform analyses with respect to the three distributions ε_0 for the stock that we introduced at the end of Section 3.2 and use preference parameters $\tau_1=1,\rho_1=3/2$ and $\tau_2=1/2,\rho_2=1$. (These risk- and skew-tolerances correspond to square-root and log utility for the first and second agents, respectively, see our later discussion in Section 5.1 and the formulas for risk- and skew-tolerances in Eq. (31).) Fig. 2 presents implied volatility for $\sigma=1/12$ when three options are traded for the uniform (first column), the normal (second column) and the lognormal distribution (third column). Row one takes K_1 = 0.1, K_2=1, K_3=1.9 and row two takes K_1=0.3, K_2=1, K_3=1.7. We see that implied volatility decreases as we increase the strike and corresponds to the implied volatility "smile" pattern that we see in equity option markets. Furthermore, note that the size of the at-the-money implied volatility (K=1) is roughly the standard deviation of the underlying security (57.7%).

Although here we illustrated pricing and implied volatility for specific risk- and skew-tolerances, we want to point out that we found qualitatively similar implied volatility for other heterogeneous agent preference specifications. In addition, note that Eq. (17) implies that the price and implied volatility of one option will not be affected by other options, see, i.e. implied volatility is robust to changes in strikes. (This robustness is very different from open interest across strikes as we will see in Section 6.2.)

Overall, pricing through the small-noise expansion technique matches qualitatively well the observed volatility "smile" in equity option markets. Pricing results derived in this section are comforting but we will not study pricing further, since our focus in this paper is on demand and open interest.

---

[7] This follows from our derivation. Note also that holdings in the skewness portfolio $X_S = V^{-1}\chi$ are of first order in σ, because V is of second and χ of third order.

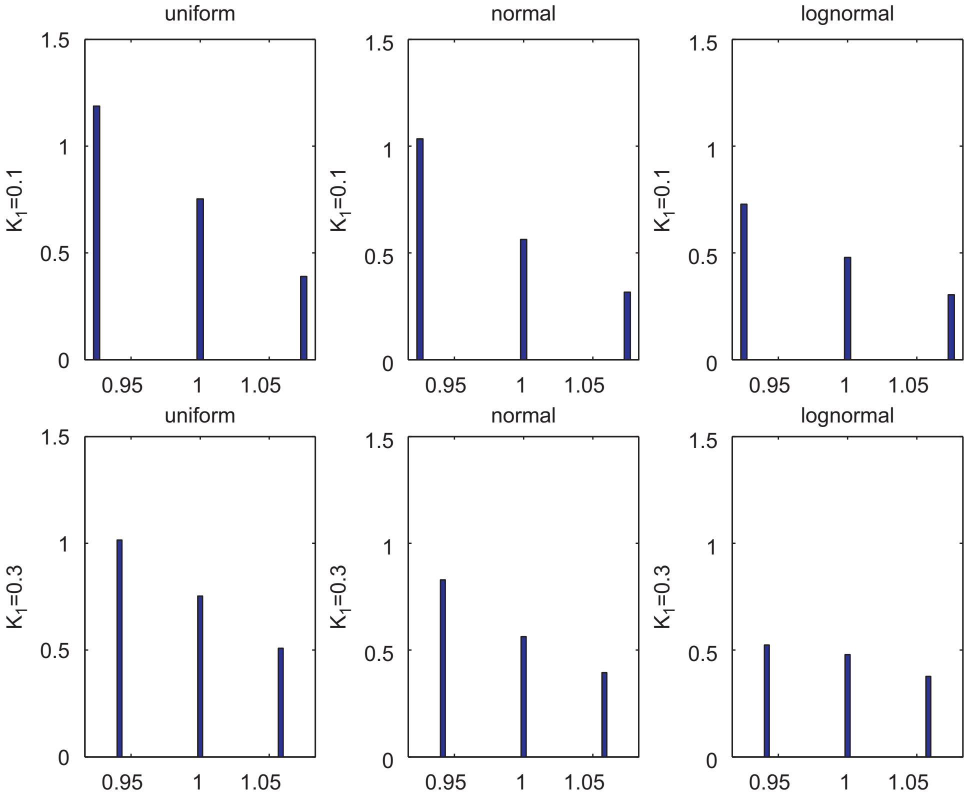

Fig. 2. The implied volatility with K_1=0.3, K_2=1, K_3=1.7 for three distributions.

4.4. Why do agents trade options?

In partial equilibrium, agents trade options, because they provide an attractive risk-return profile. Mean–variance demand will then induce agents to hold options, if risk-premiums are not set "correctly." Additional insights can be gained from general equilibrium, where risk-premiums are set endogenously to clear the market.

In general equilibrium the mean–variance tradeoff does not lead to option demand. Intuitively, this result is due to the following: mean–variance demand leads agents to hold the market portfolio. Because options are in zero net-supply, they are not contained in the market portfolio and so they are not traded. We see therefore that in general equilibrium a mean–variance analysis in a risk-sharing setup is not enough to justify option demand and that an analysis of skewness is necessary.

In general, agents have elaborate risk-preferences and would like to achieve portfolio skewness different from what the stock can offer. To see, what they strike for we need to understand the skewness portfolio: its holdings $X_{Sj} = (V^{-1}\chi)_j$ in a security j are those of a regression of the squared (de-meaned) stock payoff Y_0 on the payoff Y_j from security j; this means the skewness portfolio is the portfolio with payoff as close as possible (in terms of variance) to the squared (de-meaned) stock payoff. In short, it is the portfolio offering the payoff "closest" to the squared stock payoff.

With the interpretation that the stock represents the "market portfolio," Theorem 3 states that agents hold the market portfolio and the portfolio closest to the "squared market portfolio." Therefore we say that agents trade the market and skewness portfolios to attain the desired profile regarding market risk and the squared market risk, respectively.

Options are used to achieve their desired portfolio skewness. *If* a security with a payoff equal to the squared market would be available, *then* agents could achieve the optimal skewness level of their portfolio by trading the bond, stock and the squared stock. However, since the squared stock *cannot* be traded, agents resolve to the available securities to create a portfolio, which is as close as possible to the squared market; this portfolio is the skewness portfolio. Here, options are necessary to create the skewness portfolio, and for this reason they are traded.

5. Complete markets

In this section we assume that options can be traded for all strike parameters $0 < K < 2$ in the support of the stock distribution. In the single-period setup the market is then static complete: any payoff profile at date 1 can be attained by taking a static (buy-and-hold) position in *all* available securities, i.e. in the riskfree bond, the stock and call options with *all* strikes.

5.1. Notation

Because options are in zero net-supply, the market portfolio consists of exactly one unit of the stock and no options; this is unchanged to the incomplete markets case. However, the skewness portfolio needs to be redefined: it shall contain all options K in one unit.

A large part of our analysis does not make specific assumptions about the utility function, but for illustrative purposes we sometimes look at CPRA preferences ("constant proportional risk-aversion preferences", a.k.a. "power utility"), a class of utility functions that are often used in finance; they are characterized by a parameter $\gamma_a < 1$ and given by

$$u_a(w) = \frac{w^{\gamma_a}}{\gamma_a}. \tag{18}$$

Special cases are linear (risk-neutral) preferences with $\gamma_a = 1$, and quadratic with $\gamma_a = 2$. The limit case $\gamma_a = 0$ is to be understood as $u_a(w)=\log(w)$. A more general class of utility functions that are also often used are HARA preferences; these are characterized by the utility function

$$u_a(w) = \frac{1-\gamma_a}{\gamma_a}\left(\frac{w}{1-\gamma_a}+\beta_a\right)^{\gamma_a}, \tag{19}$$

subject to $w/(1-\gamma_a)+\beta_a \geq 0$ and $\gamma_a < 1$. This includes CPRA preferences. A special case of HARA is the negative exponential $u_a(w) = -\beta_a \exp(-\beta_a w)$ with $\gamma_a = -\infty$.

5.2. Sharing rule and option demand

It is well known, see, e.g., Duffie (2001) and Cochrane (2001), that in a complete market agents can attain any wealth profile defined on the support of the distribution. We denote by $W^*(y)$ the optimal date 1 wealth of the first agent as a function of the stock payoff $y = Y_0(\sigma)$ and refer to this as the *sharing rule*. The sharing rule relates agents' marginal utility

$$0 = \frac{\partial u_1}{\partial w}(W^*(y)) - \lambda(\sigma)\cdot\frac{\partial u_2}{\partial w}(y - W^*(y)), \tag{20}$$

where $\lambda(\sigma)$ is an intrinsic multiplier.[8]

For further analysis we need a representation of the date 0 positions through which a particular wealth profile can be attained. We use a general result derived by Bakshi and Madan (2000), Carr and Madan (2001), and Weinbaum (2009): When a riskfree bond and call options can be traded for any strike within the support of the distribution of the underlying security, then any twice continuously differentiable payoff function of the underlying security at a fixed date can be attained by taking static (buy-and-hold) positions in the riskfree bond, the underlying security and *all* the call options; the position sizes are given through derivatives of the payoff function.[9] Applying this general result to our setup, we find that agents can attain the sharing rule by a taking at date 0 a position in the riskfree bond (with size $d_{1B}(\sigma)$), the stock and in a continuum of call options; here

$$W^*(Y_S(\sigma)) = d_{1B}(\sigma) + d_{1S}(\sigma)Y_S(\sigma) + \int_0^2 d_{1K}(\sigma)\cdot(1+\varepsilon_0-K)^+\,dK, \tag{21}$$

where

$$d_{1K}(\sigma) = \sigma\cdot\frac{\partial^2 W^*}{\partial y^2}(K). \tag{22}$$

Therefore, the second derivative of the sharing rule contains all relevant information for option demand and for the remainder of this complete markets section we focus on this.

To illustrate how the sharing rule determines option demand we look at a special case where the first agent has a square-root utility and the other has log utility function (CPRA preferences with $\gamma_1 = 1/2, \gamma_2 = 0$). Then we solve Eq. (20) for the unique solution with positive wealth[10]

$$W^*(y) = \frac{1}{2\lambda^2(\sigma)}\left\{\sqrt{1+4\lambda^2(\sigma)y}-1\right\}, \tag{23}$$

[8] The sharing rule $W^*(y)$ also depends on σ through the multiplier $\lambda(\sigma)$; we dropped the dependence to simplify our exposition.

[9] The representation follows from a Taylor-series expansion of the payoff function, see, e.g. Lemma 1 and its proof in Weinbaum (2009).

[10] Wang (1996), based Dumas (1989), derived a similar closed-form expression for the sharing rule in a continuous-time setup for two heterogeneous agents with square-root and log utility.

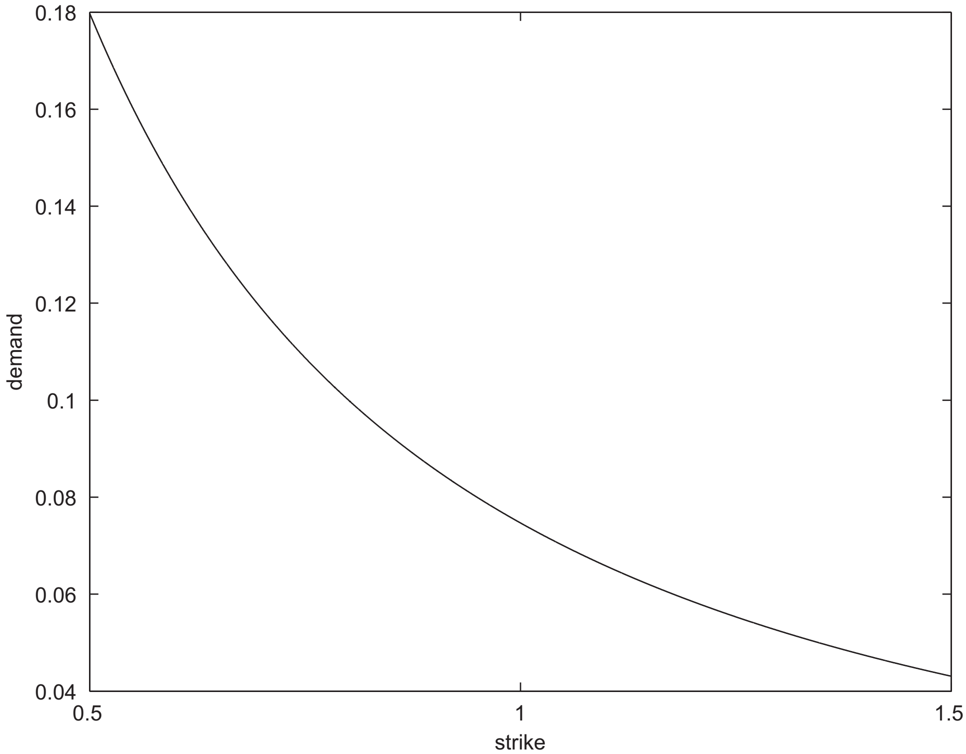

Fig. 3. The first agent's demand curve across strikes when the first agent has square-root and the second has logarithmic utility.

which implies that the first agent's demand is

$$d_{1K}(\sigma) = \sigma \frac{\partial^2 W^*}{\partial y^2}(K) = \sigma \frac{2\lambda^2(\sigma)}{(1+4\lambda^2(\sigma)K)^{3/2}}. \tag{24}$$

Note that demand is always positive, i.e. the first agent is always long the options, independent of the parameter $\lambda(\sigma)$.

It is well known that in complete markets $\theta(\partial u_1/\partial w)(W^*(Y_S(\sigma)))$ is the so-called state-price density which permits pricing all payoffs. (The constant θ is set such that the riskfree bond is priced correctly.) In particular, it permits pricing the initial endowment (one-half unit of stock) and the payoff from the sharing rule; the budget for the agent equates if both have the same price. To determine $\lambda(\sigma)$ we need to solve $E[\partial u_1/\partial w(W^*(Y_S(\sigma))) \cdot (W^*(Y_S(\sigma)) - Y_S(\sigma)/2)] = 0$; it will depend on the distribution for ε_0. To illustrate results, we assume ε_0 is the normal distribution introduced at the end of Section 3.2. We choose $\sigma = 1/2$ and solve the budget equation to get $\lambda(1/2) = 1.3963$. Fig. 3 plots the resulting demand curve. (We chose a fairly large $\sigma = 1/2$ to illustrate the demand curve over a wide range.)

5.3. *Expanding the sharing rule and the state-price density*

This section determines the expansion of the sharing rule that will be used throughout this complete markets section. To simplify our presentation[11] we assume $\lambda(\sigma)$ is constant and for all $0 \le \sigma \le 1$ equal to the value $\lambda(\sigma = 0)$, which we simply denote λ. Note that $Y_S(0)=1$, $W^*(Y_S(0))=1/2$, and so Eq. (20) implies

$$\lambda = \frac{\frac{\partial u_1}{\partial w}(1/2)}{\frac{\partial u_2}{\partial w}(1/2)}. \tag{25}$$

To determine the first and second order derivatives of wealth note that Eq. (20) holds for all y, and so its left- and right-hand side derivatives must be equal. Under the assumption that λ is constant in σ, taking derivatives repeatedly, evaluating

[11] We could study the general case via a two-dimensional Taylor-series expansion of W^* in the parameters σ *and* λ along the lines of the one-dimensional expansion of Eq. (28); note that $\lambda(\sigma)$ enters as a parameter into the sharing rule. We refrain from doing so, as it does not affect results qualitatively; it is just tedious as we need to track derivatives in λ, too. For a discussion of the error we refer to Section 7.3.

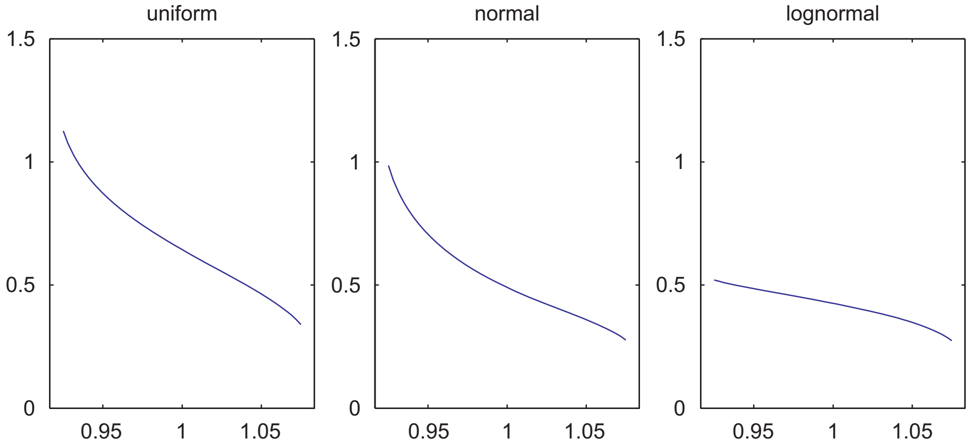

Fig. 4. The implied volatility curve across strikes for three distributions.

them at $Y_S(0)=1$ and using the equality $W^*(Y_S(0))=W^*(1)=P_S(0)/2=\frac{1}{2}$, we find that

$$\frac{\partial W^*}{\partial y}(1)=\frac{\tau_1}{\tau_1+\tau_2} \quad \text{and} \quad \frac{\partial^2 W^*}{\partial y^2}(1)=2\tau_1\tau_2\frac{\rho_1-\rho_2}{(\tau_1+\tau_2)^3}. \tag{26}$$

To illustrate pricing implications we derive the first-order expansion of the state-price density $m(Y_S(\sigma))$ in the next step. Because the interest rate is assumed to be zero, the state-price density is $m(Y_S(\sigma))=\partial u_1/\partial w(W^*(Y_S(\sigma)))/E[\partial u_1/\partial w(W^*(Y_S(\sigma)))]$. A first order series expansion of $\partial u_1/\partial w(y)$ around $Y_S(0)$, together with the above expansion of the sharing rule gives a first-order expansion of $\partial u_1/\partial w(W^*(Y_S(\sigma)))$:

$$\frac{\partial u_1}{\partial w}\left(\frac{1}{2}\right)+\frac{\partial^2 u_1}{\partial w^2}\left(\frac{1}{2}\right)\cdot\left\{\frac{\partial W^*}{\partial y}(1)\cdot(Y_S(\sigma)-1)+\frac{1}{2}\frac{\partial^2 W^*}{\partial y^2}(1)\cdot(Y_S(\sigma)-1)^2\right\}. \tag{27}$$

Dividing this by its expected value gives a first-order expansion of the state-price density. While the state-price density gives the price of any option K as $E[m(Y_S(\sigma))\cdot Y_K(\sigma)]$, this expansion of the state-price density gives us a first-order approximation of option prices.

Fig. 4 uses the same preference parametrization for both agents (square-root and logarithmic utility) that we used in Fig. 2 to study in separate plots the three stock distributions ε_0 that we introduced at the end of Section 3.2; each plot presents the implied volatility curve across strikes for $\sigma=1/12$ and any strike parameter $0.3<K<1.7$. Qualitatively, we see for all three distributions the implied volatility smile. Although we illustrated here the implied volatility curve for specific risk- and skew-tolerances, we found qualitatively similar implied volatility curves for other heterogeneous agent preference specifications.

Note that when agents are homogeneous, then $\partial^2 W^*/\partial y^2(1)=0$ and so the first-order approximation in Eq. (27) is a linear function of the stock. According to Rubinstein (1976) this implies with lognormal distributions the Black–Scholes formula for option prices. Fig. 4 suggests that heterogeneity leads to the implied volatility smile. Benninga and Mayshar (2000) showed that heterogeneity is a source of the implied volatility "smile" pattern; our expansion suggests the same. A full analysis, however, is beyond the scope of this paper.

5.4. Why do agents trade options?

While the previous analysis tells us that options are useful to attain the sharing rule, this section studies the link between option demand and risk-preferences to gain economic insights into the structure of option demand.

A straightforward approximation for $d_{aK}(\sigma)=\sigma(\partial^2 W^*/\partial y^2(K))$ of Eq. (22) is obtained by evaluating the second order derivative of the sharing rule at K=1:

$$d_{aK}^{first}(\sigma)=\sigma\cdot\frac{\partial^2 W^*}{\partial y^2}(1)=2\sigma\tau_1\tau_2\frac{\rho_1-\rho_2}{(\tau_1+\tau_2)^3}. \tag{28}$$

We refer to this as the first-order approximation of option demand, because it is of first order in σ and because it shares many features with the first-order approximation of option demand in incomplete markets. Eq. (28) links option demand to agents' skew-tolerances. It remains to clarify the link between the skewness portfolio in the complete and incomplete markets. Applying Eq. (21) analogously to the squared market (as a "fictitious" sharing rule) we find

$$\int_0^2 \sigma(1+\varepsilon_0-K)^+\,dK=\frac{Y_S^2(\sigma)}{2}, \tag{29}$$

i.e. the skewness portfolio we introduced at the beginning of this section has payoff one-half the squared stock. Therefore we can write the following first-order approximation for the overall payoff from all option positions:

$$\int_0^2 d_{aK}^{first}(\sigma)\cdot\sigma(1+\varepsilon_0-K)^+\,dK=\tau_1\tau_2\frac{\rho_1-\rho_2}{(\tau_1+\tau_2)^3}\cdot Y_S^2(\sigma). \tag{30}$$

While the riskfree bond and the stock (a.k.a. market portfolio in our setup) permit agents to achieve their optimal mean–variance portfolio position, the first order demand approximation allows agents to replicate the squared market. Note that in complete markets agents can actually create a portfolio that has the payoff of the squared market, whereas in incomplete markets they can only get a portfolio "as close as possible" to it. (Before we linked this to the skewness portfolio; this motivates the name skewness portfolio here.) We expect options to be traded as long as agents' skew-tolerances differ. This intuition carries over from the incomplete markets section.

5.5. *Who is long and who is short options for HARA preferences*

This section studies who is buying options (long position) and who is selling options (short position) when agents have HARA preferences with $\gamma_1\neq\gamma_2$. We first derive our results and then compare them with the literature.

To determine who is long and who is short options we need to sign option demand. In our expansion, this is straightforward: if $\partial^2W^*/\partial y^2(1)>0$, then the first agent's demand is positive,[12] see Eq. (28). According to Eq. (26) this requires $\rho_1>\rho_2$. We calculate for HARA preferences

$$\tau_a=\beta_a+\frac{1}{2(1-\gamma_a)},\quad \rho_a=\frac{2-\gamma_a}{2(1-\gamma_a)}. \tag{31}$$

Therefore, the first agent is long options if and only if $\gamma_1>\gamma_2$. (To see the latter, note that, ρ_a as a function of γ_a is strictly increasing.) If $\beta_1=\beta_2$ this is equivalent to $\tau_1>\tau_2$, i.e. the more risk-tolerant agent is long options. Note that the first (second) agent is then long (short) *all* traded options.

So far, only a handful papers have studied the question who is buying/selling options; to our knowledge, these are Leland (1980), Brennan and Solanki (1981), Franke et al. (1998), Dieckmann and Gallmeyer (2005), Bates (2008), and Weinbaum (2009).[13] In the following paragraphs we re-derive their results with the exception of Brennan and Solanki (1981) as it is closely related to Leland (1980) and could be derived along the same lines.

Leland (1980) compares the risk-tolerance of individual agents to that of the representative agent in complete markets. To relate our analysis to his, let us denote *in this paragraph* for the first agent $\tau_1(w)=(-\partial u_1/\partial w)/(\partial^2u_1/\partial w^2)$ and $\rho_1(w)=\frac{1}{2}\tau_1^2(w)(\partial^3u_1/\partial w^3)/(\partial u_1/\partial w)$ the risk-tolerance and skew-tolerance as functions of wealth w. (The risk-tolerance and skew-tolerance we adopt throughout the paper are these terms evaluated at one-half.) Also we denote u_* the utility function of the representative agent and by $\tau_*(w)=(-\partial u_*/\partial w)/(\partial^2u_*/\partial w^2)$ and $\rho_*(w)=\frac{1}{2}\tau_*^2(w)(\partial^3u_*/\partial w^3)/(\partial u_*/\partial w)$ her risk-tolerance and skew-tolerance as functions of wealth. Leland (1980) argues that the first agent is long options if his or risk-tolerance $\tau_1(w)$ grows faster in wealth w than the risk-tolerance of the representative agent $\tau^*(w)$, i.e. if $\partial\tau_1/\partial w>\partial u_*/\partial w$. (Otherwise, the agent is short options.) We calculate $\partial\tau_1/\partial w=2\rho_1-1$ for the first agent and similarly $\partial\tau_*/\partial w=2\rho_*-1$ for the representative agent. Therefore (with our notation) Leland (1980) states that the first agent is long options if $\rho(w)>\rho_*(w)$. Locally, around initial wealth, this will hold if $\rho(\frac{1}{2})>\rho_*(\frac{1}{2})$, which is exactly our earlier result.

In a continuous-time model with jumps, Bates (2008) considers heterogeneous agents with preferences over terminal wealth *and* the number of crashes, while Dieckmann and Gallmeyer (2005) study heterogeneous agents with preferences over wealth. Both find that the more risk-averse agent is short options. Weinbaum (2009) finds the same result, see his proposition 2. Our expansion developed the same intuition for call options.

Another interesting paper that addressed option demand is by Franke et al. (1998). They take a setup where two agents (both with the *same* HARA preferences) face a marketable risk on which they can trade a complete set of options; they also face an uninsurable so-called background risk. Due to the modeling of background risk, the setup of Franke et al. (1998) does *not* fit into our setup; but our setup can be easily adjusted to provide an analysis.

Appendix B.3 extends our small-noise expansion to an analysis of sharing rules under background risk. We provide the mathematical details about the derivation for Franke et al. (1998) in that Appendix and only sketch the analysis here: The sharing rules, risk- and skew-tolerance parameters need to be modified slightly; we denote these here $\overline{W}^*,\overline{\tau},\overline{\rho}$. With these parameters replacing the ones we use throughout this paper, our analysis remains valid. In particular, in our small-noise expansion version of Franke et al. (1998), the first agent is long the options if

$$\frac{\partial^2\overline{W}^*}{\partial y^2}(1)=\overline{\tau}_1\overline{\tau}_2\frac{\overline{\rho}_1-\overline{\rho}_2}{(\overline{\tau}_1+\overline{\tau}_2)^3}>0,$$

[12] This is only the first order demand approximation. However, it determines the sign of demand, see our discussion of errors in Section 7.3.

[13] Wang (1996), based on an analysis by Dumas (1989), did not look at option demand, but derived a sharing rule similar to ours in Eq. (23) which we used to determine demand. There we noted that the first agent is always long options. The preferences looked at were HARA with $\gamma_1=1/2$ and $\gamma_2=0$. Since $\gamma_1>\gamma_2$, our expansion suggest that the first (second) agent is long (short) all options, in line with their implication.

i.e. if $\overline{\rho}_1 > \overline{\rho}_2$. A very illustrative example in Franke et al. (1998) is when only the second agent faces background risk; they find that the first agent then has a concave sharing rule, i.e. sells options on the market portfolio. Here, with our expansion we find the same result, see Appendix B.3.

6. Open interest curve across strikes

In this section we first present a stylized fact about the shape of the open interest curve. Then we analyze what that shape would be in (general) equilibrium for incomplete and complete markets.

6.1. The stylized fact

Several plain vanilla options with various exercise prices and expiration months are available for trade at any date on the same underlying security. In this section, we look at the dependence of open interest on the strike price for fixed maturity and underlying security. Option volume and open interest have been documented, among others, by Lakonishok et al. (2007), Buraschi and Jiltsov (2006), and Bollen and Whaley (2004). It is a stylized fact that a plot of that dependence for calls and puts shows a peak near to the at-the-money contract; this feature appears whatever the underlying stock, index or maturity.

As an example we will present that feature for the options on Microsoft (ticker symbol: `MSFT`) and the NASDAQ 100 tracking unit (ticker symbol: `QQQQ`) with March 2006 maturity.[14] Expiration for those series was on Saturday March 18, 2006. The option series with the shortest maturity is the most actively traded one and therefore we take a snapshot on the last four Fridays preceding expiry.

The structure of the market is the following: March 2006 options on `MSFT` can be traded with strike prices in \$2.5 intervals between \$7.5 and \$45 and those on `QQQQ` with strike prices \$1 intervals between \$24 and \$52. During that period contracts were also traded with maturity at the "Saturday immediately following the third Friday" in April 2006, July 2006, October 2006, January 2007, and January 2008 for `MSFT`, and in April 2006, May 2006, June 2006, September 2006, January 2007, and January 2008 for `QQQQ`.

The left column of Fig. 5 presents `MSFT` and the right column `QQQQ`; rows one to four look separately at Friday February 24, March 3, March 10 and March 17, 2006. The closing prices for `MSFT` (`QQQQ`) on those four Fridays were \$26.63 (\$41.26), \$26.93 (\$41.45), \$27.17 (\$40.56), and \$27.50 (\$41.45), respectively. Each plot shows these through a vertical line. On all four Fridays and for both underlying securities Fig. 5 shows the stylized fact that there is a peak of open interest near to the at-the-money call option. Note that the stylized fact is evident whatever maturity we look at: it holds at very short horizons of one day up to three weeks.[15]

6.2. Theoretical shape of the open interest curve in incomplete markets

This section studies the shape of the open interest curve that results from our first-order approximation of demand when the number of options is finite and "small." (If the number of options is not "small," the market would be "almost complete" and the analysis of the complete markets Section 6.3 should be used to analyze the shape of the open interest curve.)

We noted in Section 4 that equilibrium demand in options will arise from skew-tolerance and co-skewness as long as $\rho_1 \neq \rho_2$ and that all agents hold the riskfree bond, the market and the skewness portfolios. Then, option demand is determined entirely by the multiplicative product of preference parameters $\tau_1\tau_2(\rho_1-\rho_2)/(\tau_1+\tau_2)^3$) and the composition of the skewness portfolio $X_S = V^{-1}\chi$. To analyze the shape of the open interest curve, it is therefore sufficient to analyze the absolute composition of the skewness portfolio $|X_S|$.

The composition of the skewness portfolio depends on the location of the strike grid and the distribution, only. We perform analyses with respect to the three distributions ε_0 for the stock that we introduced at the end of Section 3.2: Fig. 6 illustrates the composition of the skewness portfolio – the shape of the open interest curve – for[16] $\sigma = 1/12$ when three options are traded for the uniform (first column), the normal (second column) and the lognormal distribution (third column). Row one takes $K_1 = 0.1$, $K_2=1$, $K_3=1.9$ and row two takes $K_1=0.3$, $K_2=1$, $K_3=1.7$. (This choice of strike parameters has been studied previously in Fig. 2 where we illustrated the implied volatility across strikes in incomplete markets. We found there the implied volatility "smile" pattern and this pattern was robust to changes in strikes.)

Fig. 6 shows in the first row a "dip" for the open interest curve, whatever the distribution; this shape is clearly not in line with the stylized fact. However, looking at other strike grids, we see that this (negative) result is not robust: in the second row, i.e. for slight changes in the strikes, we see in the left-most plot that a "peak" for the open interest curve

[14] Consistently, options on `MSFT` and `QQQQ` are among the most actively traded ones written on stocks and indices.

[15] Open interest is negligible for options with strikes outside the depicted range. The grid is identical for calls and puts, i.e. for every call with a maturity T and a strike K also a put with same maturity and strike can be traded. Put-call parity allows replicating a long put position in a strike K through joint long positions in the stock, the call with strike K and an investment of \$$K$ in the bond. Translating all positions into corresponding call positions and creating a new summary open interest in a call with strike K by adding to it the open interest in the put with strike K yields similar pictures.

[16] For illustrative purposes we need to fix σ, but this does not impact the relative size of the skewness portfolio components, i.e. the *shape* of the open interest curve.

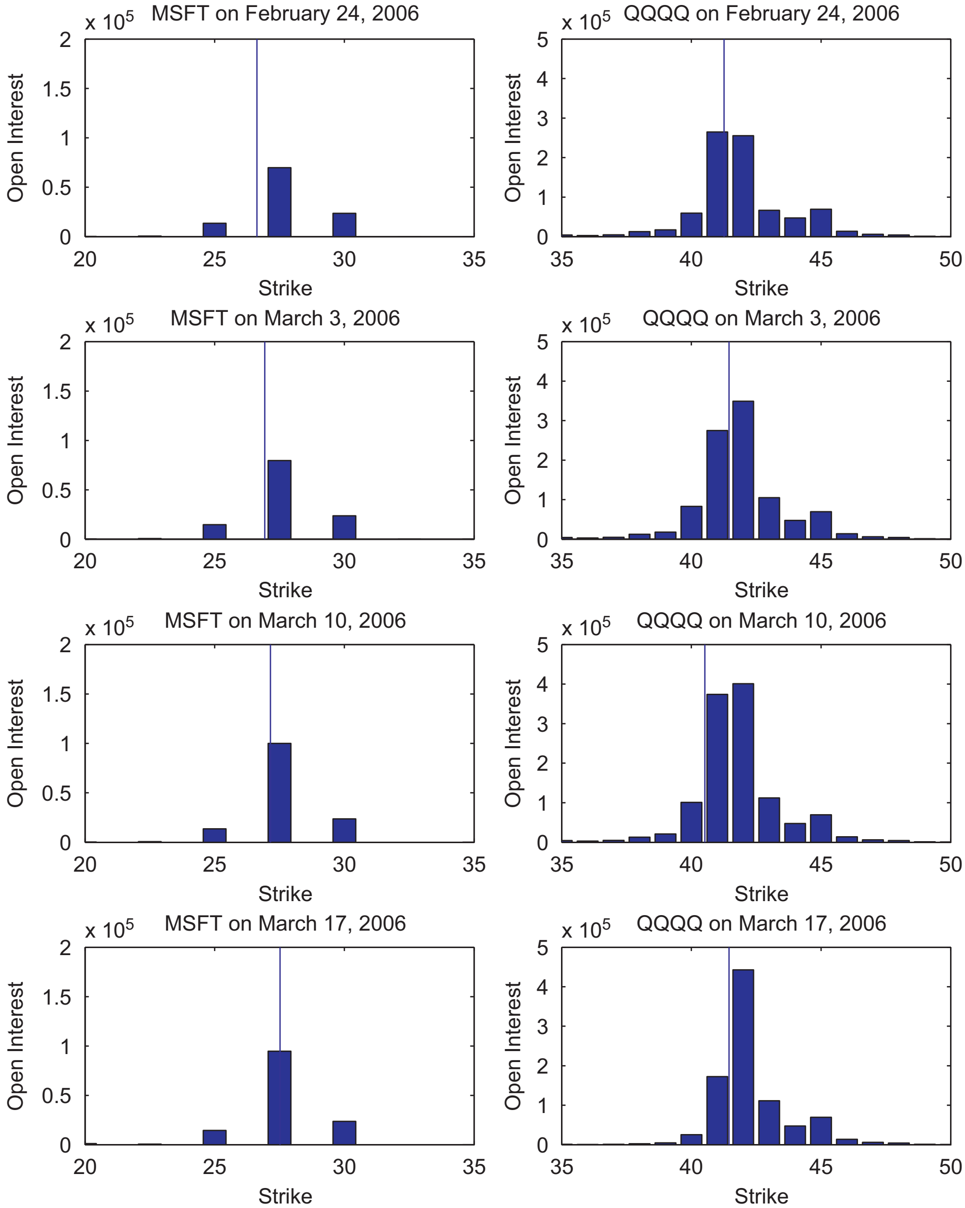

Fig. 5. Empirically observed call option open interest in the March 2006 series for Microsoft (MSFT) and the NASDAQ 100 tracking unit (QQQQ) on the last four Friday's before maturity. (The vertical line indicates the price of the underlying on that day.)

appears. Yet, this peak is not very pronounced and it is surprising that the sign of the curvature for the open interest curve flips by slight changes in the strikes. We see similar results in columns two and three; there, the shape also changes considerably in response to change in the strikes: in the second column the shape changes from a "dip" to an "almost flat" curve while in the third column it changes from a "dip" to a monotonically decreasing curve. Overall, we sometimes see a peak and sometimes a dip, and the presence of a peak/dip is sensitive to the strike grid and to the distribution used. Further plots we looked at confirmed this for other strike grids and also showed that the shape of the open interest curve is sensitive to the number of traded options.

The sensitivity of the skewness portfolio $X_S = V^{-1}\chi$ should not be surprising to the reader familiar with mean–variance portfolio construction. In mean–variance theory, agents' trade off the portfolio mean against its variance; it is well known that portfolio holdings are very sensitive to the input parameters (means and variances), see, e.g. Black and Litterman (1992). In the construction of the skewness portfolio we trade off portfolio skewness against its variance. By analogy we can therefore expect that the composition of the skewness portfolio and therefore the shape of the open interest curve are

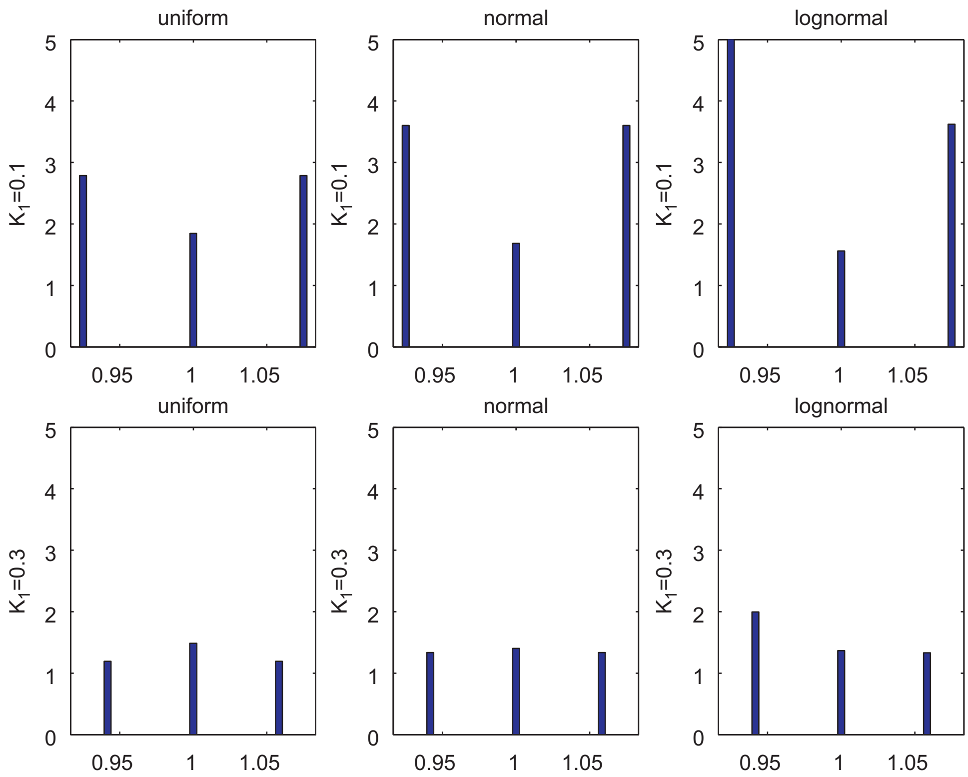

Fig. 6. The shape of the open interest curve across strikes for two strike grids with three traded options each and three distributions.

very sensitive to the input parameters (location of the strike grid, number of options, type of distribution). This is what we see in our plots for the shape of the open interest curve.

We conclude that in incomplete markets with a small number of traded options the shape of the open interest curve is very sensitive to the number of options traded, to the type of distribution and the location of the strike grid. This is in sharp contrast to the robustness of the stylized fact.

6.3. *Theoretical shape of the open interest curve in complete markets*

This section studies the shape of the open interest curve when options are traded for all strikes in (0,2), so that the market is complete.

In Section 5.2 we looked at a case where the two agents have square-root and log utility, respectively, and plotted in Fig. 3 the first agent's demand curve for $\sigma = 1/2$ with the actual $\lambda(\sigma)$ when the stock distribution ε_0 is a truncated normal. Because demand is always positive, Fig. 3 shows the open interest curve. The curve is monotonically decreasing and of convex shape; it clearly does *not support the stylized fact*. This *qualitative* result is independent of assumptions about the stock distribution. To see this, note that the actual distribution only enters via the $\lambda(\sigma)$ which solves the budget equation. However, the resulting demand Eq. (24) is for *all* values of $\lambda(\sigma)$ always positive, montononically decreasing and convex.

One might think that this is due to the risk-preferences we used. Therefore, we studied the sharing rule in closed-form for several other CPRA preferences: the sharing rule $W^*(y)$ is by Eq. (20) the solution to

$$W^*(y) - (\lambda(\sigma))^{1/(\gamma_1-1)} \cdot (y - W^*(y))^{(\gamma_2-1)/(\gamma_1-1)} = 0, \tag{32}$$

and when the ratio $(\gamma_2-1)/(\gamma_1-1)$ is an integer, this equation describes a polynomial. Their solutions are often known in closed-form, e.g. the symbolic algebra of the MATHEMATICA program can provide them for $\gamma_2 = 0$ and $(\gamma_2-1)/(\gamma_1-1) = 2,3,4,5,\ldots,10$ (corresponding to $\gamma_1 = 1/2, 2/3, 3/4, 4/5, \ldots, 9/10$). We checked the resulting sharing rules, however, in all of them the theoretical open interest curve is strictly decreasing and convex. *None* of the cases we looked at *supported the stylized fact*. Overall, this analysis suggests that a risk-sharing setup with CPRA preferences cannot justify the shape of the open interest curve we see in the markets.

The sharing rule and thereby the open interest curve cannot be determined in closed-form for general (non-CPRA) preferences. However, we can easily carry out the expansion for all given functional forms of the utility function to determine the shape of the open interest curve. The first-order approximation of Section 5.4 allowed us to gain economic insights into the structure of option demand. In the incomplete markets case the first-order approximation determined the shape of the open interest curve, see Fig. 6: It was *not* flat across strikes and exhibited curvature. In the complete markets case of this section, however, the first order demand is flat across strikes, see Eq. (28), so to determine the *shape* we need to look at a higher-order approximation. A Taylor series expansion of $\sigma(\partial^2 W^*/\partial y^2(K))$ around $y=1$ gives the following third-order approximation of option demand:

$$d_{aK}^{third}(\sigma) = \sigma \cdot \frac{\partial^2 W^*}{\partial y^2}(1) + (K-1) \cdot \sigma \cdot \frac{\partial^3 W^*}{\partial y^3}(1) + \frac{1}{2}(K-1)^2 \cdot \sigma \cdot \frac{\partial^4 W^*}{\partial y^4}(1). \tag{33}$$

This describes $d_{aK}(\sigma)$ as a quadratic polynomial in the strike K. The expansion terms have interesting interpretations: the first, second and third terms in this equation determine the level, the slope and the curvature of option demand, respectively. Here, we see now that when $\partial^2 W^*/\partial y^2(1)$ is positive, demand $d_{aK}(\sigma)$ is also positive and the sign of the curvature term $\partial^4 W^*/\partial y^4(1)$ will determine if the open interest curve shows a peak or a dip. (Appendix B.1 writes these terms out as second-, third- and fourth-order derivatives of agents' utility functions, each evaluated at wealth of one-half.)

For the remainder of this section we assume that agents have HARA preferences with $1 > \gamma_1 > \gamma_2$ to analyze if the open interest curve peaks. This class includes CPRA preferences, but HARA preferences have two additional degrees of freedom (β_1 and β_2). For small σ, the first agent will be long in options and the open interest will be equal to that agent's demand. Therefore, we will have a peak whenever $\partial^4 W^*/\partial y^4(1)$ is strictly negative, and a dip whenever it is strictly positive.

The $\partial^4 W^*/\partial y^4(1)$ term can be determined in closed-form by symbolic algebra, but the terms are fractional polynomials in $\beta_1,\beta_2,\gamma_1,\gamma_2$ and hard to interpret. We refrain here from presenting them, because we are mainly interested in its (negative) minimal value ("maximal peak" of the open interest curve). We asked the MATHEMATICA program to determine the minimum of $\partial^4 W^*/\partial y^4(1)$ for all parameter choices $\beta_1,\beta_2,\gamma_1,\gamma_2$ under the constraints that the first agent is long in options ($\gamma_1 > \gamma_2$), that $\gamma_1,\gamma_2 < 1$ and that the positivity restriction $1/2/(1-\gamma_a)+\beta_a \geq 0$ holds. (The positivity constraint after (19) is evaluated here at $w=1/2$.) The result is that the minimum of $\partial^4 W^*/\partial y^4(1)$ is attained at $\beta_1 = -2.8933$, $\beta_2 = 8.9604$, $\gamma_1 = 0.9164$, $\gamma_2 = -11.6036$. With these parameters $\partial^2 W^*/\partial y^2(1) = 0.1869$, $\partial^3 W^*/\partial y^3(1) = 0.0418$, and $\partial^4 W^*/\partial y^4(1) = -0.1044$. The latter number is negative; this means that the open interest curve *peaks*.

Unfortunately, the number $\partial^4 W^*/\partial y^4(1) = -0.1044$ is only slightly negative, so the peak is not very pronounced. To see this better we plot in Fig. 7 the resulting open interest for $\sigma = 1/2$, based on Eq. (33). (We chose a fairly large σ to illustrate the shape over a wide range.) We see a peak but the peak is not very pronounced. Clearly, even for the parameters with maximal curvature this does not support the observed peak in the open interest curve. The overall conclusion is that with HARA preferences the resulting open interest curve does not match the stylized fact.

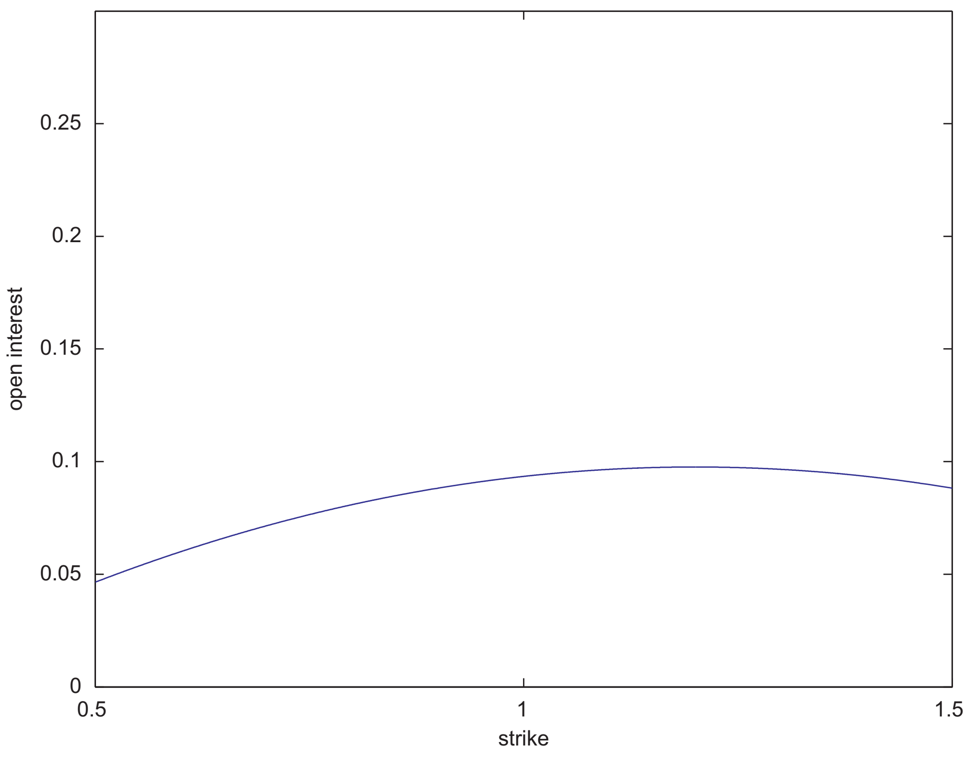

Fig. 7. The complete markets open interest curve with maximal negative curvature for two agents with HARA preferences.

7. Robustness of our analysis

Our paper has derived a theory of option demand and open interest by looking at the leading terms in an expansion of demand and risk-premiums in a simplified setup. Throughout this paper we assumed agents only derived utility from wealth and traded at date 0 while payoffs occurred at date 1. Furthermore, we assumed the economy only contains two agents. The following two sections we explain how our expansion could be extended to cover the general setup. Yet, these expansions come at the expense of additional complexity, while they do not add economic insights, do not affect the shape of the open interest curve, nor do they affect the size of demand considerably. The last Section 7.3 addresses the errors that result from approximations we made.

7.1. Two dates and no consumption

We could easily generalize our two date setup to multiple trading periods. In complete markets, note that agents only would have to trade today to achieve the sharing rule (their optimal payoff profile) at any future date, so our analysis already covers that case.

In incomplete markets, the Bellman principle of dynamic optimization permits us to break the multi-period maximization problem of each agent into a series of single-period problems using the so-called indirect utility function. While we looked only at call options throughout this paper, our closed-form approximations do not depend on this: we presented first-order demand as a function of mean, variance and skewness of the date 1 payoff in partial and general equilibrium but did not use the specific structural payoff profile of call options. The formulas still apply with the interpretation that "future payoffs" are next trading date's prices. We could define risk- and skew-tolerances for the indirect utility function; because these are evaluated at wealth in the "no-risk" economy, they boil down to risk- and skew-tolerances of agents' utility functions, evaluated at the wealth 1/2 in the no-risk economy. So the resulting preference parameters could still be treated as numbers "known" at date 0.

We could therefore proceed in a backward induction to determine demand and risk-premiums (prices). Risk-premiums will then depend on the stock price in previous periods and therefore we will see intertemporal hedge demand in addition to the risk-sharing demand we looked at in this paper. While such an analysis is feasible and the modeling is more realistic, it does not add qualitative insights, nor does it affect the size of demand and the shape of the open interest curve significantly. Therefore, we refrain from carrying out this analysis here.

We could also generalize our small-noise expansion to incorporate consumption. The demand for risky assets will then be determined by the marginal utility of future consumption; this could be captured by our analysis of future wealth in conjunction with the multi-period approach described above. A consumption-based analysis would permit us to endogenize the interest rate, but it would come at the expense of additional complexity and bring us away from our focus on option demand. For simplicity we therefore ignore consumption and set the interest rate equal to 0.

7.2. Multiple agents

For incomplete markets, Appendix A.4 derives a version of Theorems 2 and 3 with multiple agents: the three-fund separation of Theorem 3 will continue to hold with multiple agents; only that the reference point of an agent will no longer be the "other" agent but a so-called "average" agent. In consequence, the actual size of each agent's demand for the skewness portfolio will change, but the insights remain valid with multiple agents, that agents trade options to achieve the squared market return, the relative size of option demand for each agent will be unchanged and the shape of the open interest curve continues to be determined by the relative composition of the skewness portfolio.

Our complete markets analysis is based on Eq. (20), i.e. that an agent's marginal utility is proportional to the other's marginal utility, across all states of wealth; this was the starting point for a series expansion of marginal utility to infer structural properties of wealth and option demand. With multiple agents, agent's marginal utility is proportional to each other's and so our analysis could be carried out analogously to the one presented there. Summing up long positions of all agents in an option we could determine the option's open interest. Looking at multiple agents will not affect our results qualitatively, because the general insight that agents trade options to achieve the desired squared stock payoff remains unchanged. For a mathematical type discussion of how to carry out the generalization to multiple agents we refer to Section B.2 in the Appendix.

7.3. Errors of the small-noise expansion

Our small-noise expansion leads to an approximation $d_{aK}^{first}(\sigma)$ of option demand. Addressing the resulting error is the main purpose of this section. Furthermore we fixed the value of $\lambda(\sigma)$ in Section 5.4 for all values of σ at $\lambda(0) = \partial u_1/\partial w(1/2)/\partial u_2/\partial w(1/2)$, see Eq. (25); we denote here $d_{aK}^{\lambda}(\sigma)$ the resulting demand approximation and address

Table 1
Maximal relative errors in option demand.

	Uniform (%)	Normal (%)	Lognormal (%)
$e^{\lambda}(\sigma)$			
$\sigma=1/52$	0.0008	0.0008	0.0008
$\sigma=1/12$	0.0343	0.0232	0.0194
$\sigma=2/12$	0.1462	0.0965	0.0793
$e^{first}(\sigma)$			
$\sigma=1/52$	2.5739	2.5739	2.5739
$\sigma=1/12$	11.2751	11.2875	11.2916
$\sigma=2/12$	22.8438	22.9049	22.9261
$e^{third}(\sigma)$			
$\sigma=1/52$	0.0019	0.0019	0.0019
$\sigma=1/12$	0.1182	0.1078	0.1044
$\sigma=2/12$	0.7948	0.7525	0.7378

its error. We analyze maximal relative errors across all options $0<K<2$ as

$$e^{first}(\sigma)=\max_{0<K<2}\left|\frac{d_{aK}^{first}(\sigma)-d_{aK}(\sigma)}{d_{aK}(\sigma)}\right|,\quad e^{\lambda}(\sigma)=\max_{0<K<2}\left|\frac{d_{aK}^{\lambda}(\sigma)-d_{aK}(\sigma)}{d_{aK}(\sigma)}\right|.$$

The reference case for $d_{aK}(\sigma)$ is here the closed-form expressions we derived in the complete markets case of two agents with square-root and log utility (Section 5.2). Table 1 presents maximal errors for all three distributions introduced in Section 3.2 for three choices of $\sigma=1/52,1/12,2/12$. (These correspond to the horizons of one week, one month and two months, respectively, according to our discussion in Section 3.2.)

We see that errors are larger when σ is greater; this is what we expect from a series expansion. The time-lengths we study are within the range used in applications, with the typical one being less than one month, i.e. $\sigma\leq 1/12$. Our approximation of $\lambda(0)$ shows errors $e^{\lambda}(\sigma)$ that are less than 0.15%; for our typical application of one month the error is less than 0.04%. These errors are small and show that ignoring higher order terms in $\lambda(\sigma)$ does not affect the size of demand significantly. Similarly, when we look at first-order approximations $e^{first}(\sigma)$ we see that with the two month horizon errors are less than 23.0% and with the time frame of one month we have in mind, it is less than 11.3%. While these numbers affect demand they are in our view still fairly precise and permit us to infer that the main demand terms are skewness driven. (For informative purposes we also present third order demand errors $e^{third}(\sigma)$, which we define similar to first order demand errors $e^{first}(\sigma)$ using the third order demand approximation $d_{aK}^{third}(\sigma)$ instead of the first order demand approximation $d_{aK}^{first}(\sigma)$ of Eq. (33). We see that the accuracy increases considerably.)

The errors of the first order demand approximation and of the $\lambda(0)$ approximation are too small to affect who is holding the long and the short position and will also not affect the shape considerably. In particular the errors are too small to change our insights from Section 6 that the risk-sharing setup has great difficulties to justify the stylized fact about the shape of the open interest curve.

8. Conclusion

This paper looked at a sequence of risk-sharing economies to analyze equilibrium demand and open interest in options. We studied the first-order demand and risk-premium approximations in closed-form to gain insights into option demand. We pointed out that in partial equilibrium securities are priced by their marginal contribution to portfolio and squared portfolio risk; mean–variance theory may lead to option demand if risk-premiums are set "incorrectly." In the market-clearing general equilibrium mean–variance option demand disappears and we need to study preferences over skewness. We then argued that agents trade options to achieve a portfolio payoff as close as possible to the squared stock payoff; agents hold three funds: a riskfree bond, the stock (the surrogate for the "market" portfolio) and the so-called skewness portfolio. We illustrated the resulting implied volatilities across strikes; they match qualitatively the observed implied volatility "smile" pattern in equity markets. Finally, we confronted the risk-sharing theory with a stylized fact that the curve of open interest across strikes peaks close to the at-the-money option. In complete markets (with HARA preferences) and in incomplete markets we saw that a risk-sharing setup does not support the stylized fact of open interest.

Appendix A. Incomplete markets analysis

We prove in this Appendix Theorem 1 for agent a=1,2. Technical problems arise; their root is that in the $\sigma=0$ economy all securities coincide and therefore their demand is indeterminate. For this we make use of the perturbation analysis of Judd and Guu (2001). Our derivation is based on ε_0 and de-meaned options' payoffs $\varepsilon_j=(1+\varepsilon_0-K_j)^+-E[(1+\varepsilon_0-K_j)^+]$ instead of securities' payoffs $Y(\sigma)$. Therefore, here we define the (new) variance–covariance matrix $\tilde{V}$ and (co)-skewness vector $\tilde{\chi}$ by

$$\tilde{V}_{jk} = E[\varepsilon_j \varepsilon_k], \quad \tilde{\chi}_j = E[\varepsilon_0^2 \varepsilon_j],$$

for agent a=1,2, $\tilde{\zeta}_a$ the N+1 dimensional *(third order) co-moment vector* $\tilde{\zeta}_a = Cov(\varepsilon,(\varepsilon^{\perp} \cdot d_a(0))^2)$ and the vector of holdings $\tilde{X}_S = \tilde{V}^{-1}\tilde{\chi}$. Note that $V = \tilde{V} \cdot \sigma^2$ and $X_S = \tilde{X}_S \cdot \sigma$. As in Section 4 we refer to the stock as security number 0.

A.1. Preparing for the implicit function theorem

In each σ–economy total wealth of agent a=1,2 at date 1 is

$$W_{1a}(\sigma) = W_{0a}(\sigma) + \sum_{j=0}^{N} d_{aj}(\sigma) \cdot (Y_j(\sigma) - P_j(\sigma)). \tag{A.1}$$

For each agent a=1,2 we define the (N+1)-dimensional function $H_a(d_a(\sigma),\sigma)$ by

$$H_{aj}(d_a(\sigma),\sigma) = \frac{1}{\sigma}\frac{\partial E[u_a(W_{1a})]}{\partial d_{aj}} = E\left[\frac{\partial u_a}{\partial w}(W_{1a}(\sigma)) \cdot (\varepsilon_j + \pi_j(\sigma)\sigma)\right], \tag{A.2}$$

for j=0,1,…,N. The *first-order conditions*[17] of agent a=1,2 in the σ economy are then $H_a(d_a(\sigma),\sigma) = 0$.

Our idea is to derive agent's holdings through an application of the implicit function theorem: taking σ–derivatives of H_a, we get (purely formally) $0 = \partial H_a/\partial d_a \cdot \partial d_a/\partial\sigma + \partial H_a/\partial\sigma$, i.e. $\partial d_a/\partial\sigma$ as the ratio of $\partial H_a/\partial\sigma$ and $\partial H_a/\partial d_a$. Yet, the way our expansion is set up, we see below that $\partial H_a/\partial d_a(d_a(0),\sigma = 0) = 0$: in the $\sigma = 0$ economy all securities coincide and therefore their demand is indeterminate. In the spirit of Hospital's rule we "require" $\partial H_a/\partial\sigma(d_a(0),\sigma = 0)$ to be also equal to zero and use second order derivatives of H_a to express $\partial d_a/\partial\sigma$. Formally this requires a generalized implicit function theorem, introduced by Judd and Guu (2001):

Theorem 5 (*Theorem 7 in* Judd and Guu (2001)). *Suppose $H : R^n \times R \rightarrow R^n$ is analytic, and $H(x,0) = 0$ for all $x \in R^n$. Furthermore, suppose that for some $(x_0,0)$*

$$\frac{\partial H}{\partial x}(x_0,0) = 0_{n\times n}, \frac{\partial H}{\partial \sigma}(x_0,0) = 0_n, \quad \text{and det} \left(\frac{\partial^2 H}{\partial x \partial \sigma}(x_0,0)\right) \neq 0.$$

Then there is an open neighborhood $\mathcal{N}$ of $(x_0,0)$, and a function $h(\sigma) : R \rightarrow R^n, h(\sigma) \neq 0$ for $\sigma \neq 0$, such that

$$H(h(\sigma),\sigma) = 0 \quad \textit{for } (h(\sigma),\sigma) \in \mathcal{N}.$$

Furthermore, h is analytic and can be approximated by a Taylor series. In particular, the first-order derivatives equal

$$\frac{\partial h}{\partial \sigma}(0) = -\frac{1}{2} \cdot \left(\frac{\partial^2 H}{\partial x \partial \sigma}(\sigma)\right)^{-1} \cdot \frac{\partial^2 H}{\partial \sigma^2}(\sigma).$$

Here we will simply apply their procedure. Since $W_{0a}(\sigma) = \frac{1}{2}P_0(\sigma) = \frac{1}{2}(1-\pi_0(\sigma)\sigma^2)$ the derivatives of terminal wealth are, treating $d_a(\sigma)$ as an implicit function,

$$\frac{\partial W_{1a}}{\partial \sigma}(\sigma) = -\frac{1}{2}\left(\frac{\partial \pi_0}{\partial \sigma}(\sigma)\sigma^2 + \pi_0(\sigma)2\sigma\right) + \sum_{j=0}^{N} d_{aj}(\sigma) \cdot \left(\varepsilon_j + \frac{\partial \pi_j}{\partial \sigma}(\sigma)\sigma^2 + \frac{\partial \pi_j}{\partial \sigma}(\sigma)2\sigma\right),$$

$$\frac{\partial^2 W_{1a}}{\partial \sigma^2}(\sigma) = -\frac{1}{2}\left(\frac{\partial^2 \pi_0}{\partial \sigma^2}(\sigma)\sigma^2 + \frac{\partial \pi_0}{\partial \sigma}(\sigma)4\sigma + \pi_0(\sigma)2\right) + \sum_{j=0}^{N} d_{aj}(\sigma) \cdot \left(\frac{\partial^2 \pi_j}{\partial \sigma^2}(\sigma)\sigma^2 + \frac{\partial \pi_j}{\partial \sigma}(\sigma)4\sigma + \frac{\partial \pi_j}{\partial \sigma}(\sigma)2\right).$$

Note that $W_{1a}(0)$ = 1/2, that $E[\partial W_{1a}/\partial\sigma(0)] = 0$, that $\partial^2 W_{1a}/\partial\sigma^2(0) = -\pi_0(0) + \sum_{j=0}^{N} d_{aj}(0)2\pi_j(0)$ and that $\tilde{\zeta}_{aj} = E[(\partial^2 W_{1a}/\partial\sigma^2(0))^2 \cdot \varepsilon_j]$.

A.2. Calculating derivatives

To apply Theorem 5 we will now calculate the first-order derivatives of H_a with respect to d and σ

$$\frac{\partial H_{aj}}{\partial d_{ak}}(\sigma) = E\left[\frac{\partial^2 u_a}{\partial w^2}(W_{1a}(\sigma)) \cdot (\varepsilon_j + \pi_j(\sigma)\sigma) \cdot (\varepsilon_k \sigma + \pi_k(\sigma)\sigma^2)\right],$$

$$\frac{\partial H_{aj}}{\partial \sigma}(\sigma) = E\left[\frac{\partial^2 u_a}{\partial w^2}(W_{1a}(\sigma)) \cdot \frac{\partial W_{1a}}{\partial \sigma}(\sigma) \cdot (\varepsilon_j + \pi_j(\sigma)\sigma)\right] + E\left[\frac{\partial u_a}{\partial w}(W_{1a}(\sigma)) \cdot \left(\pi_j(\sigma) + \frac{\partial \pi_j}{\partial \sigma}(\sigma)\sigma\right)\right].$$

[17] To simplify the exposition we divided by σ when defining H_a in Eq. (A.2).

We also calculate the necessary second order derivatives of H_a with respect to d and σ

$$\frac{\partial^2 H_{aj}}{\partial d_{ak}\partial\sigma}(\sigma)=E\Big[\frac{\partial^3 u_a}{\partial w^3}(W_{1a}(\sigma))\cdot\frac{\partial W_{1a}}{\partial\sigma}(\sigma)\cdot(\varepsilon_j+\pi_j(\sigma)\sigma)(\varepsilon_k\sigma+\pi_k(\sigma)\sigma^2)$$
$$+\frac{\partial^2 u_a}{\partial w^2}(W_{1a}(\sigma))\cdot\left(\pi_j(\sigma)+\frac{\partial\pi_j}{\partial\sigma}(\sigma)\sigma\right)\cdot(\varepsilon_k\sigma+\pi_k(\sigma)\sigma^2)$$
$$+\frac{\partial^2 u_a}{\partial w^2}(W_{1a}(\sigma))\cdot(\varepsilon_j+\pi_j(\sigma)\sigma^2)\cdot\left(\varepsilon_k+\frac{\partial\pi_k}{\partial\sigma}(\sigma)\sigma^2+\pi_k(\sigma)2\sigma\right)\Big],$$

$$\frac{\partial^2 H_{aj}}{\partial\sigma^2}(\sigma)=E\left[\frac{\partial^3 u_a}{\partial w^3}(W_{1a}(\sigma))\cdot\left(\frac{\partial W_{1a}}{\partial\sigma}(\sigma)\right)^2\cdot(\varepsilon_j+\pi_j(\sigma)\sigma)+\frac{\partial^2 u_a}{\partial w^2}(W_{1a}(\sigma))\cdot\frac{\partial W_{1a}}{\partial\sigma}(\sigma)\cdot(\varepsilon_j+\pi_j(\sigma)\sigma)\right.$$
$$\left.+2\frac{\partial^2 u_a}{\partial w^2}(W_{1a}(\sigma))\cdot\frac{\partial W_{1a}}{\partial\sigma}(\sigma)\cdot\left(\pi_j(\sigma)+\frac{\partial\pi_j}{\partial\sigma}(\sigma)\sigma\right)+\frac{\partial u_a}{\partial w}(W_{1a}(\sigma))\cdot\left(2\frac{\partial\pi_j}{\partial\sigma}(\sigma)+\frac{\partial^2\pi_j}{\partial\sigma^2}(\sigma)\sigma\right)\right].$$

A.3. Deriving partial equilibrium demand

We saw in Appendix A.2 that $\partial H_a/\partial d(\sigma=0)=0$ and that $\partial^2 H_a/\partial d_a\partial\sigma(\sigma=0)=\partial^2 u_a/\partial w^2(\frac{1}{2})\cdot\tilde{V}$. The utility function u_a is strictly concave and $\tilde{V}$ invertible, so that $\partial^2 H_a/\partial d_a\partial\sigma(\sigma=0)$ is invertible with

$$\left(\frac{\partial H_a^2}{\partial d_a\partial\sigma}(0)\right)^{-1}=\frac{1}{\frac{\partial^2 u_a}{\partial w^2}\left(\frac{1}{2}\right)}\tilde{V}^{-1}.$$

To apply Theorem 5 we require that $H_\sigma(\sigma=0)=0$, i.e.

$$0=\frac{\partial H_{aj}}{\partial\sigma}(0)=\frac{\partial^2 u_a}{\partial w^2}\left(\frac{1}{2}\right)\cdot\sum_{k=0}^{N}d_{ak}(0)Cov(\varepsilon_j,\varepsilon_k)+\frac{\partial u_a}{\partial w}\left(\frac{1}{2}\right)\cdot\pi_j(0),$$

which means that

$$0=\frac{\partial^2 u_a}{\partial w^2}\left(\frac{1}{2}\right)\cdot\tilde{V}\cdot d_a+\frac{\partial u_a}{\partial w}\left(\frac{1}{2}\right)\cdot\pi(0),\text{ i.e. } d_a(0)=\tau_a\cdot(\tilde{V}^{-1}\pi(0)).$$

This describes the zero order demand term. Now that the conditions have been checked to apply Theorem 5 we calculate that

$$\frac{\partial^2 H_{aj}}{\partial\sigma^2}(0)=\frac{\partial^3 u_a}{\partial w^3}\left(\frac{1}{2}\right)\cdot E\left[\left(\frac{\partial W_{1a}}{\partial\sigma}(0)\right)^2\varepsilon_j\right]+2\cdot\frac{\partial^2 u_a}{\partial w^2}\left(\frac{1}{2}\right)\cdot E\left[\frac{\partial W_{1a}}{\partial\sigma}(0)\right]\pi_j(0)+2\cdot\frac{\partial u_a}{\partial w}\left(\frac{1}{2}\right)\cdot\frac{\partial\pi_j}{\partial\sigma}(0),\text{ i.e.}$$

$$\frac{\partial^2 H_{aj}}{\partial\sigma^2}(0)=\frac{\partial^3 u_a}{\partial w^3}\left(\frac{1}{2}\right)\cdot\tilde{\zeta}_j+2\cdot\frac{\partial u_a}{\partial w}\left(\frac{1}{2}\right)\cdot\frac{\partial\pi_j}{\partial\sigma}(0).$$

Theorem 5 tells us then that

$$\frac{\partial d_a}{\partial\sigma}(0)=-\frac{1}{2}\left(\frac{\partial^2 H_a}{\partial d_a\partial\sigma}(\sigma=0)\right)^{-1}\frac{\partial^2 H_a}{\partial\sigma^2}=\frac{\rho_a}{\tau_a}(\tilde{V}^{-1}\tilde{\zeta}_a)+\tau_a\left(\tilde{V}^{-1}\frac{\partial\pi}{\partial\sigma}(0)\right).$$

A.4. Deriving the general equilibrium allocation

Although our analysis in the paper looks only at the two-agent case, this section calculates the equilibrium for an economy populated by a=1,…,A agents with given risk- and skew-tolerance parameters τ_a,ρ_a: Individual demand is assumed to be described by the demand equations derived for the two-agent case, i.e. by Eq. (14). We leave the aggregate supply of the stock and the option unchanged to the two-agent case, i.e. one and zero, respectively.

For reference in this subsection, we define the aggregate risk-tolerance $\overline{\tau}$ and the aggregate skew-tolerance $\overline{\rho}$ by

$$\overline{\tau}=\sum_{a=1}^{A}\tau_a,\overline{\rho}=\frac{1}{\overline{\tau}}\sum_{a=1}^{A}\tau_a\rho_a.$$

Using $d_a(0)=\tau_a\cdot\tilde{V}^{-1}\pi(0)$ we get that zero order aggregate demand is $\overline{\tau}\cdot\tilde{V}^{-1}\pi(0)$. This must equal the market portfolio vector X_M, which implies

$$\pi(0)=\frac{1}{\overline{\tau}}(\tilde{V}\cdot X_M)=\frac{1}{\overline{\tau}}\tilde{V}_{0j}\quad\text{or equivalently}\quad\pi(0)=\beta\cdot\pi_0(0),\tag{A.3}$$

with β describing the vector with $\beta_j = \tilde{V}_{0j}/\tilde{V}_{00}$. This implies

$$d_a(0) = \frac{\tau_a}{\overline{\tau}} X_M. \tag{A.4}$$

From this we derive $\tilde{\zeta}_a = Cov(\varepsilon,(\varepsilon^{\perp} \cdot d_a(0))^2) = \tau_a^2/\overline{\tau}^2 \cdot Cov(\varepsilon,\varepsilon_0^2) = \tau_a^2/\overline{\tau}^2 \cdot \tilde{\chi}_j$. Therefore

$$\frac{\partial d_a}{\partial \sigma}(0) = \tilde{V}^{-1}\left(\tau_a \frac{\partial \pi}{\partial \sigma}(0) + \frac{\rho_a \tau_a}{\overline{\tau}^2}\tilde{\chi}\right).$$

The market-clearing condition for the first-order demand is then

$$0 = \sum_{a=1}^{A} \frac{\partial d_a}{\partial \sigma}(0) = \tilde{V}^{-1}\left(\overline{\tau}\frac{\partial \pi}{\partial \sigma}(0) + \frac{\overline{\rho}}{\overline{\tau}}\tilde{\chi}\right).$$

Therefore

$$\frac{\partial \pi}{\partial \sigma}(0) = -\frac{\overline{\rho}}{\overline{\tau}^2}\tilde{\chi}, \quad \text{or equivalently } \frac{\partial \pi}{\partial \sigma}(0) = \gamma_j \cdot \frac{\partial \pi_0}{\partial \sigma}(0),$$

with γ describing the vector with $\gamma_j = \tilde{\chi}_j/\tilde{\chi}_0$. This implies

$$\frac{\partial d_a}{\partial \sigma}(0) = (\rho_a - \overline{\rho})\frac{\tau_a}{\overline{\tau}^2}(\tilde{V}^{-1}\tilde{\chi}).$$

Note that in the two-agent case $\overline{\tau} = \tau_1 + \tau_2$, which implies the first-order premium coincides of Theorem 2. Furthermore in the two-agent case $\overline{\rho} = (\rho_1\tau_1 + \rho_2\tau_1)/(\tau_1+\tau_2)$, so that

$$(\rho_1 - \overline{\rho})\frac{\tau_1}{\overline{\tau}^2} = \left(\frac{\rho_1\tau_1}{(\tau_1+\tau_2)^2} - \frac{\tau_1}{(\tau_1+\tau_2)^3}(\rho\tau_1 + \rho_2\tau_2)\right) = -\tau_1\tau_2\frac{\rho_1-\rho_2}{(\tau_1+\tau_2)^3},$$

which coincides with the first-order demand term of Theorem 2 since $X_S = (\tilde{V}^{-1}\tilde{\chi})\sigma$.

Appendix B. Complete markets analysis

B.1. Higher-order results

In addition to the risk- and skew-tolerance defined in Eq. (12) we define for each agent a=1,2 the *kurtosis-tolerance* κ_a and *pentosis-tolerance* ξ_a by

$$\kappa_a = \frac{\tau_a^3}{3}\frac{\frac{\partial^4 u_a}{\partial w^4}\left(\frac{1}{2}\right)}{\frac{\partial u_a}{\partial w}\left(\frac{1}{2}\right)}, \quad \xi_a = \frac{\tau_a^4}{4}\frac{\frac{\partial^5 u_a}{\partial w^5}\left(\frac{1}{2}\right)}{\frac{\partial u_a}{\partial w}\left(\frac{1}{2}\right)}. \tag{B.5}$$

As before, these terms are ratios of derivatives of agent's utility functions, evaluated at one-half, which we treat as (non-random) numbers. Whereas risk- and skew-tolerances incorporate first-, second- and third-order derivatives of agents' utility functions, kurtosis- and pentosis-tolerance add fourth and fifth-order derivatives of agents' utility functions. (Kurtosis-tolerance is related to the risk-vulnerability concept of Gollier and Pratt, 1996; to our knowledge, pentosis-tolerance has not been studied in the literature before.)

By continuing the procedure leading to Eq. (26) we derive

$$\frac{\partial^3 W^*}{\partial y^3}(1) = 3\tau_1\tau_2\frac{\kappa_1-\kappa_2}{(\tau_1+\tau_2)^4} + 12\tau_1\tau_2(\tau_1\rho_2+\tau_2\rho_1)\frac{\rho_1-\rho_2}{(\tau_1+\tau_2)^5}, \tag{B.6}$$

$$\begin{aligned}\frac{\partial^4 W^*}{\partial y^4}(1) &= 4\tau_1\tau_2\frac{\xi_1-\xi_2}{(\tau_1+\tau_2)^5} + 12\tau_1\tau_2\frac{\kappa_2(-5\rho_2\tau_1+\rho_1(3\tau_1-2\tau_2))-\kappa_1(-5\rho_1\tau_2+\rho_2(3\tau_2-2\tau_1))}{(\tau_1+\tau_2)^6} \\ &\quad + 24\tau_1\tau_2(\rho_1-\rho_2)\frac{5\rho_2^2\tau_1^2+\rho_1^2\tau_2^2+\rho_1\rho_2(\tau_1^2-8\tau_1\tau_2+\tau_2^2)}{(\tau_1+\tau_2)^7}.\end{aligned} \tag{B.7}$$

For HARA utility we calculate

$$\kappa_a = -\frac{(3-\gamma_a)(2-\gamma_a)}{3(1-\gamma_a)^2} \quad \text{and} \quad \xi_a = \frac{(4-\gamma_a)(3-\gamma_a)(2-\gamma_a)}{4(1-\gamma_a)^3}. \tag{B.8}$$

B.2. The multi-agent setup

Throughout, this paper only looked at the two-agent case, but could easily be generalized to the multi-agent case (a=1, …, A agents with $A>2$). To see how, we denote here W_a^* the optimal date 1 wealth profile of agent a. Then, marginal utility equates across agents, i.e. Eq. (20) holds similarly for all combinations $a=2,\ldots,A: 0=\partial u_1/\partial w(W_1^*(y))-\lambda_{1a}(\sigma)\cdot\partial u_a/\partial w(W_a^*(y))$, for suitable $\lambda_{1a}(\sigma)$. Open interest in option K is the positive part of all agents' demand $\sum_{a=1}^{A}(d_{aK}(\sigma))^+=\sigma\sum_{a=1}^{A}(\partial^2 W_a^*/\partial y^2(K))^+$. This means we would need to do a structural analysis of all wealth profiles $W_a^*(y)$. In Section 6.3 we performed a structural analysis of the sharing rule in the two agent case, only, to analyze the shape of the open interest curve. Summing up positions across more than two agents will not affect our results qualitatively. To simplify our exposition we only looked at the simpler two-agent setup throughout.

B.3. Franke–Stapleton–Subrahmanyam

Our setup can be easily modified to provide an analysis of Franke et al. (1998). They take two agents with identical HARA utility functions ($u_1(w)=u_2(w)=u(w)$, facing an uninsurable background risk. We note that the defining equation for the sharing rule $\overline{W}^*(y)$, previously in Eq. (20), now reads (with our simplification that λ does not depend on σ)

$$0=E\left[\frac{\partial u}{\partial w}(\overline{W}^*(y)+e_1)\middle|y\right]-\lambda\cdot E\left[\frac{\partial u}{\partial w}(y-\overline{W}^*(y)+e_2)\middle|y\right], \tag{B.9}$$

where e_1,e_2 is a non-positive random variable describing the background risk of the first and second agent. The analogy to our analysis is as follows: Eq. (B.9) means all terms in derivatives of u need to be replaced by expected values, with expectation taken over the background risk, i.e. *here*

$$\overline{\tau}_a=-\frac{E\left[\frac{\partial u}{\partial w}\left(\frac{1}{2}+e_a\right)\right]}{E\left[\frac{\partial^2 u}{\partial w^2}\left(\frac{1}{2}+e_a\right)\right]},\quad \overline{\rho}_a=\frac{\overline{\tau}_a^2}{2}\frac{E\left[\frac{\partial^3 u}{\partial w^3}\left(\frac{1}{2}+e_a\right)\right]}{E\left[\frac{\partial u}{\partial w}\left(\frac{1}{2}+e_a\right)\right]}.$$

Our analysis of Section 5.5 then remains valid, in particular the sign of $\partial^2\overline{W}^*/\partial y^2(1)=\overline{\tau}_1\overline{\tau}_2(\overline{\rho}_1-\overline{\rho}_2/(\overline{\tau_1}+\overline{\tau}_2)^3)$ determines whether the first agent is long or short. In particular, the first agent is long options if $\overline{\rho}_1>\overline{\rho}_2$.

A very illustrative example in Franke et al. (1998) is when only the second agent faces background risk ($e_1=0$); they find that the first agent then has a concave sharing rule, i.e. sells options on the market portfolio. This means in our expansion that $\overline{\tau}_1,\overline{\rho}_1$ are the usual ones and $\overline{\tau}_2,\overline{\rho}_2$ are as above. To determine if $\overline{\rho}_1>\overline{\rho}_2$ note first that in this case

$$\overline{\rho}_1=\frac{\frac{\partial u}{\partial w}\left(\frac{1}{2}\right)\cdot\frac{\partial^3 u}{\partial w^3}\left(\frac{1}{2}\right)}{\left(\frac{\partial^2 u}{\partial w^2}\left(\frac{1}{2}\right)\right)^2}\quad\text{and}\quad \overline{\rho}_2=\frac{E\left[\frac{\partial u}{\partial w}\left(\frac{1}{2}+e_2\right)\right]\cdot E\left[\frac{\partial^3 u}{\partial w^3}\left(\frac{1}{2}+e_2\right)\right]}{\left(E\left[\frac{\partial^2 u}{\partial w^2}\left(\frac{1}{2}+e_2\right)\right]\right)^2}.$$

It is known for HARA utility that $\partial u/\partial w,\partial^3 u/\partial w^3$ are convex and $\partial^2 u/\partial w^2$ is concave. Therefore, by Jensen's inequality

$$E\left[\frac{\partial u}{\partial w}\left(\frac{1}{2}+e_2\right)\right]<\frac{\partial u}{\partial w}\left(\frac{1}{2}+E[e_2]\right),$$

$$E\left[\frac{\partial^3 u}{\partial w^3}\left(\frac{1}{2}+e_2\right)\right]<\frac{\partial^3 u}{\partial w^3}\left(\frac{1}{2}+E[e_2]\right),$$

$$E\left[\frac{\partial^2 u}{\partial w^2}\left(\frac{1}{2}+e_2\right)\right]>\frac{\partial^2 u}{\partial w^2}\left(\frac{1}{2}+E[e_2]\right).$$

This implies with Eq. (31) that

$$\overline{\rho}_1=\frac{2-\gamma_1}{2(1-\gamma_1)}=\frac{\frac{\partial u}{\partial w}\left(\frac{1}{2}+E[e_2]\right)\cdot\frac{\partial^3 u}{\partial w^3}\left(\frac{1}{2}+E[e_2]\right)}{\left(\frac{\partial^2 u}{\partial w^2}\left(\frac{1}{2}+E[e_2]\right)\right)^2}<\frac{E\left[\frac{\partial u}{\partial w}\left(\frac{1}{2}+e_2\right)\right]E\left[\frac{\partial^3 u}{\partial w^3}\left(\frac{1}{2}+e_2\right)\right]}{\left(E\left[\frac{\partial^2 u}{\partial w^2}\left(\frac{1}{2}+e_2\right)\right]\right)^2}=\rho_2.$$

This means that according to our expansion the first agent is short options. This result is identical with the result of Franke et al. (1998).

References

Ait-Sahalia, Y., Brandt, M., 2001. Variable selection for portfolio choice. The Journal of Finance 56, 1297–1351.

Arrow, K.J., 1953. Le rôle des valeurs boursières pour la répartition la meilleure des risques. In: Econométrie, Colloques Internationaux du Centre National de la Recherche Scientifique, pp. 41–47.

Back, K., 1993. Asymmetric information and options. The Review of Financial Studies 6, 432–472.

Bakshi, G., Kapadia, N., Madan, D., 2001. Stock return characteristics, skew laws, and the differential pricing of individual equity options. Review of Financial Studies 16 (1), 101–143.

Bakshi, G., Madan, D., 2000. Spanning and derivative security valuation. Journal of Financial Economics 55 (2), 101–143.

Bates, D.S., 2008. The market for crash risk. Journal of Economic Dynamics and Control 32 (7), 2291–2321.
Benninga, S., Mayshar, J., 2000. Heterogeneity and option pricing. Review of Derivatives Research 4, 7–27.
Bhamra, H., Uppal, R., 2009. The effect of introducing a non-redundant derivative on the volatility of stock-market returns when agents differ in risk aversion. Review of Financial Studies 22, 2303–2330.
Black, F., Litterman, R., 1992. Global portfolio optimization. Financial Analysts Journal 48, 28–43.
Bollen, N.P., Whaley, R., 2004. Does net buying pressure affect the shape of implied volatility functions? The Journal of Finance 59, 711–753
Brandt, M., 1999. Estimating portfolio and consumption choice: a conditional Euler equations approach. The Journal of Finance 54, 1609–1646.
Brennan, M.J., Cao, H.H., 1996. Information, trade, and derivative securities. The Review of Financial Studies 9, 163–208.
Brennan, M.J., Solanki, R., 1981. Optimal portfolio insurance. Journal of Financial and Quantitative Analysis 16, 279–300.
Buraschi, A., Jiltsov, A., 2006. Model uncertainty and option markets with heterogeneous beliefs. The Journal of Finance 61 (6), 2841–2897.
Carr, P., Madan, D., 2001. Optimal positioning in derivative securities. Quantitative Finance 1, 19–37.
Chang, Y.P., Johnson, H., Schill, M.J., 2006. Asset pricing when returns are nonnormal: Fama–French factors vs. higher order systematic co-moments. Journal of Business 79 (2), 923–940.
Cochrane, J.H., 2001. Asset Pricing. Princeton University Press.
Debreu, G., 1959. Theory of Value. Yale University Press.
Detemple, J., Selden, L., 1991. A general equilibrium analysis of option and stock market interactions. International Economic Review 32, 279–303.
Dieckmann, S., Gallmeyer, M.F., 2005. The equilibrium allocation of diffusive and jump risks with heterogeneous agents. Journal of Economic Dynamics and Control 29 (9), 1547–1576.
Dittmar, R., 2002. Nonlinear pricing kernels, kurtosis preference, and the cross-section of equity returns. The Journal of Finance 57 (1), 369–403.
Driessen, J., Maenhout, P., 2007. An empirical perspective on option pricing anomalies. Review of Finance 11, 561–603.
Duffie, D., 2001. Dynamic Asset Pricing Theory, third ed. Princeton University Press.
Dumas, B., 1989. Two-person dynamic equilibrium in the capital market. The Review of Financial Studies 2 (2), 157–188.
Easley, D., O'Hara, M., Srinivas, P., 1998. Option volume and stock prices: evidence on where informed traders trade. The Journal of Finance 53 (2), 431–465.
Franke, G., Stapleton, R.C., Subrahmanyam, M.G., 1998. Who buys and who sells options: the role of options in an economy with backround risk. Journal of Economic Theory 82, 89–109.
Garleanu, N., Pedersen, L.H., Poteshman, A., 2009. Demand-based option pricing. The Review of Financial Studies 22 (10), 4259–4299.
Gollier, C., Pratt, J.W., 1996. Risk vulnerability and the tempering effect of background risk. Econometrica 64, 1109–1124.
Harvey, C.R., Siddique, A., 2000. Conditional skewness in asset pricing tests. The Journal of Finance 55 (3), 1263–1295.
Haugh, M.B., Lo, A.W., 2001. Asset allocation and derivatives. Quantitative Finance 1, 45–72.
Huang, C.F., Litzenberger, R.H., 1988. Foundations for Financial Economics. Prentice-Hall.
John, K., Koticha, A., Subrahmanyam, M., 1992. The micro-structure of options markets: informed trading, liquidity, volatility and efficiency. NYU Working Paper.
Judd, K.L., Guu, S.-M., 2001. Asymptotic methods for asset market equilibrium analysis. Economic Theory 18, 127–157.
Kimball, M., 1990. Precautionary saving in the small and in the large. Econometrica 58 (1), 53–73.
Kogan, L., Uppal, R., 2001. Risk aversion and optimal portfolio policies in partial and general equilibrium economies. MIT Working Paper.
Kraus, A., Litzenberger, R.H., 1976. Skewness preference and the valuation of risk assets. The Journal of Finance 31 (4), 1085–1100.
Kraus, A., Smith, M., 1996. Heterogeneous beliefs and the effect of replicatable options of asset prices. The Review of Financial Studies 9, 723–756.
Lakonishok, J., Lee, I., Pearson, N.D., Poteshman, A.M., 2007. Option market activity. The Review of Financial Studies 20, 813–857.
Leland, H.E., 1980. Who should buy portfolio insurance? The Journal of Finance 35, 581–594
Liu, J., Pan, J., 2003. Dynamic derivative strategies. Journal of Financial Economics 69, 401–430.
Pan, J., Poteshman, A., 2006. The information in option volume for future stock prices. The Review of Financial Studies 19, 871–908.
Ross, S.A., 1976. Options and efficiency. The Quarterly Journal of Economics 90, 74–89.
Rubinstein, M., 1976. The valuation of uncertain income streams and the pricing of options. The Bell Journal of Economics 7, 407–425.
Samuelson, P.A., 1970. The fundamental approximation theorem of portfolio analysis in terms of means, variances and higher moments. Review of Economic Studies 37, 537–542.
Vanden, J., 2006. Option coskewness and capital asset pricing. The Review of Financial Studies 19, 1279–1320.
Wang, J., 1996. The term structure of interest rates in a pure exchange economy with heterogeneous investors. Journal of Financial Economics 41 (1), 75–110.
Weinbaum, D., 2009. Investor heterogeneity asset pricing and volatility dynamics. Journal of Economic Dynamics and Control 33 (7), 1379–1397.